arXiv:1601.03524v2 [math.ST] 10 Jun 2016

DEGREES OF FREEDOM FOR PIECEWISE LIPSCHITZ ESTIMATORS

FREDERIK RIIS MIKKELSEN AND NIELS RICHARD HANSEN

Abstract. A representation of the degrees of freedom akin to Stein's lemma is given for a class of estimators of a mean value parameter in $\mathbb{R}^n$. Contrary to previous results our representation holds for a range of discontinues estimators. It shows that even though the discontinuities form a Lebesgue null set, they cannot be ignored when computing degrees of freedom. Estimators with discontinuities arise naturally in regression if data driven variable selection is used. Two such examples, namely best subset selection and debiased lasso, are considered in detail in this paper. For debiased lasso the general representation leads to an estimate of the degrees of freedom based on the lasso solution path, which in turn can be used for estimating the risk of debiased lasso. A similar estimate is proposed for best subset selection. The usefulness of the risk estimates for selecting the number of variables is demonstrated via simulations with a particular focus on debiased lasso.

1. Introduction

Representations of the effective dimension of a statistical model have been studied extensively in many different frameworks. For classical model selection criteria such as AIC and Mallows's C_p the dimension of the parameter space is used to adjust the empirical risk for its optimism so as to provide a fair model score across different dimensions. A number of extensions to models or methods without a well defined dimension exist, such as the trace of the smoother matrix for scatter plot smoothers, see e.g. Hastie & Tibshirani (1990), and the use of the divergence of a sufficiently differentiable estimator based on Stein's lemma as described in Efron (2004). Stein's lemma was used by Zou et al. (2007) and Tibshirani & Taylor (2012) to demonstrate that for the lasso estimator in a linear regression model with Gaussian errors, the number of estimated non-zero parameters is an appropriate estimate of the effective dimension.

It is well known that neither Mallows's C_p nor AIC or related information criteria correctly adjust for the optimism that results from selecting one model among a number of models of equal dimension. The usage of such methods for model selection without adequate adjustments was called "a quiet scandal in the statistical community" by Breiman (1992), who proposed a bootstrap based method for risk estimation as an alternative. Ye (1998) defined the notion of generalized degrees of freedom for an estimator of the mean in a Gaussian model and showed how to use this number for risk estimation. The results by Ye apply to discontinuous

2010 *Mathematics Subject Classification.* 62J05, 62J07.

Key words and phrases. best subset selection, debiased lasso, degrees of freedom, Stein's Lemma.

estimators that involve model selection, but his proposal for computing the degrees of freedom was similarly to Breiman's based on refitting models to perturbed data.

If the estimator satisfies the differentiability requirements for Stein's lemma, Lemma 2 in Stein (1981), the divergence of the estimator w.r.t. the data is an unbiased estimate of the degrees of freedom in the generalized sense of Ye (1998). This was used by Meyer & Woodroofe (2000), Zou et al. (2007), Kato (2009) and Tibshirani & Taylor (2012) among others to derive formulas for the degrees of freedom of estimators that are Lipschitz continuous.

For estimators with discontinuities Stein's lemma generally breaks down and the divergence will not be an unbiased estimate of the degrees of freedom. Note that an estimator can be continuous or even differentiable almost everywhere – it can be a projection locally – and still be defined globally in such a way that it has non-ignorable discontinuities. This is, in particular, the case in regression when data adaptive variable selection is used to select among a number of projection estimators. Best subset selection is one central example, but variable selection procedures lead in general to non-ignorable discontinuities. A variable selection procedure effectively divides the sample space into a finite number of disjoint regions, with the estimator being a projection, say, on each region. The resulting estimator consisting of a selection step and a projection step will generally be discontinuous on the boundary between two regions.

Estimators defined as solutions of non-convex optimization problems may have non-ignorable discontinuties as well. In regression, non-convexity can arise from nonlinear parametrizations of the mean value or from the use of data adaptive or non-convex penalties. Examples of the latter include SCAD, MCP and k-norms for $k \in [0, 1)$. The non-convexity of the penalties is closely related to their abilities to perform variable selection. Penalization with the 0-norm yields, for instance, the Lagrangian formulation of best subset selection.

Tibshirani (2015) recently made headway with the computation of the degrees of freedom for some discontinuous estimators. Specifically, he considered a linear regression model with an orthogonal design and showed how to compute the degrees of freedom for hard thresholding, which for orthogonal designs is equivalent to the Lagrangian formulation of best subset selection. He also gave an extension of Stein's lemma to some discontinuous estimators, though it was not shown if this extension applies to subset selection estimators. Hansen & Sokol (2014) gave a different generalization of Stein's lemma for all estimators that are metric projections onto a closed set. This generalization applies to subset selection and other estimators with non-convex constraints, but did not lead to a readily computable representation of the contribution to the degrees of freedom that are due to the discontinuities of the metric projection.

The main result in this paper is Theorem 2.4, which is a version of Stein's lemma for estimators that are locally Lipschitz continuous on each of a finite number of open sets, whose union makes up almost all of $\mathbb{R}^n$. This is a general class containing a number of regression estimators that include variable selection. As a main example the debiased lasso estimator in a linear regression setup – closely related to the LARS–OLS hybrid in Efron et al. (2004) and a limit case of the relaxed lasso as considered in Meinshausen (2007) – is investigated in detail in Section 3. For the debiased lasso an estimate of the degrees of freedom is derived, which only involves the computation of the lasso solution path. A simulation study

reported in Section 4 demonstrated that the resulting risk estimate leads to reliable model selection across a range of different designs and parameter settings.

For the Lagrangian formulation of best subset selection the general representation of degrees of freedom holds, but the situation is more complicated than for debiased lasso. However, it is possible to derive an approximation, which is exact for orthogonal designs, as shown in Section 5. The proof of Theorem 2.4 and some auxiliary technical results and arguments are in the appendix.

2. A general representation of degrees of freedom

Throughout the paper we consider the multivariate normal model $\mathcal{N}(\xi, \sigma^2 I)$ on $\mathbb{R}^n$ with ξ the unknown parameter, and we let $\hat{\xi} : \mathbb{R}^n \to \mathbb{R}^n$ denote an estimator of ξ. A typical application is to linear regression estimators of the form $X\hat{\beta}$ where X denotes an $n \times p$ matrix and $\hat{\beta}$ denotes an estimator of the parameters in the linear regression model. When the estimator $\hat{\beta}$ sets some of the parameters to exactly zero we say that the estimator does variable selection. The lasso, Tibshirani (1996), is an example of a globally Lipschitz continuous estimator that does variable selection, while best subset selection is a discontinuous estimator that does variable selection. The debiased lasso – as studied intensively in Section 3 – is another example of a discontinuous regression estimator that does variable selection. Though discontinuous regression estimators that do variable selection constitute the main motivation for the present paper, the general results are more conveniently formulated in terms of estimators of the mean ξ without reference to the regression setup.

Letting $Y \sim \mathcal{N}(\xi, \sigma^2 I)$ the risk of the estimator is defined as

$$\text{Risk}(\hat{\xi}) \coloneqq E\|\xi - \hat{\xi}(Y)\|_2^2,$$

provided that $\hat{\xi}(Y)$ has finite second moment, which will thus be assumed throughout. The risk is a quantification of the error of $\hat{\xi}$, hence a good estimate of this quantity is desirable. Our main interest is to estimate the risk under the normal model. The following definition introduces two notions of degrees of freedom that are useful when we want to estimate the risk. In the definition, $\psi(y; \xi, \sigma^2)$ denotes the density for the $\mathcal{N}(\xi, \sigma^2 I)$ distribution and $\langle \cdot, \cdot \rangle$ denotes the standard inner product on $\mathbb{R}^n$. The divergence operator is also needed. It is the differential operator defined as

$$\text{div}(f) = \sum_{i=1}^n \partial_i f_i$$

for $f : \mathbb{R}^n \to \mathbb{R}^n$ Lebesgue almost everywhere differentiable and with ∂_i denoting the partial derivative w.r.t. the ith coordinate.

Definition 2.1. *For a measurable map $\hat{\xi} : \mathbb{R}^n \to \mathbb{R}^n$ such that $\hat{\xi}(Y)$ has finite second moment the degrees of freedom of $\hat{\xi}$ is*

$$\text{df}(\hat{\xi}) \coloneqq \sum_{i=1}^n \frac{\text{cov}(Y_i, \hat{\xi}(Y)_i)}{\sigma^2} = \int \frac{\langle y - \xi, \hat{\xi}(y) \rangle}{\sigma^2} \psi(y; \xi, \sigma^2) dy. \tag{1}$$

If $\hat{\xi}$ is differentiable in Lebesgue almost all points and $\text{div}(\hat{\xi})$ has finite first moment Stein's degrees of freedom of $\hat{\xi}$ is

$$\text{df}_{\text{S}}(\hat{\xi}) \coloneqq E(\text{div}(\hat{\xi}(Y))). \tag{2}$$

A simple expansion of the risk yields

$$\text{Risk} = E\|Y - \hat{\xi}(Y)\|_2^2 - n\sigma^2 + 2\sigma^2 \mathrm{df}(\hat{\xi}). \tag{3}$$

Hence $\|Y - \hat{\xi}(Y)\|_2^2 - n\sigma^2 + 2\sigma^2 \widehat{\mathrm{df}}$ is an unbiased risk estimate if $\widehat{\mathrm{df}}$ is an unbiased estimate of $\mathrm{df}(\hat{\xi})$. In practice, σ^2 must be estimated as well and a bias of $\widehat{\mathrm{df}}$ can also be preferable if it reduces the variance. Hence exact unbiasedness of a risk estimate based on (3) is of secondary interest, but it is of interest to find adequate corrections of the squared error $\|Y - \hat{\xi}(Y)\|_2^2$ that can be used for model assessment and comparison.

If $\hat{\xi}$ is *almost differentiable* then $\mathrm{df}(\hat{\xi}) = \mathrm{df}_{\mathrm{S}}(\hat{\xi})$ due to Stein's lemma (Lemma 2 in Stein (1981)), in which case $\mathrm{div}(\hat{\xi})(Y)$ is an unbiased estimate of $\mathrm{df}(\hat{\xi})$. However, most estimators with discontinuities are not almost differentiable and for such estimators it is not clear if $\mathrm{div}(\hat{\xi})(Y)$ is a useful estimate of the degrees of freedom. Indeed, our main result, Theorem 2.4, provides a representation of $\mathrm{df}(\hat{\xi}) - \mathrm{df}_{\mathrm{S}}(\hat{\xi})$, which is nonzero for a range of estimators. The result provides the theoretical basis for establishing more adequate estimates of the degrees of freedom and thus the risk. Furthermore, Theorem 3.2 provides a quite remarkable connection between $\mathrm{df}(\hat{\xi})$ and $\mathrm{df}_{\mathrm{S}}(\hat{\xi})$ for the debiased lasso estimator, which can be used to derive an estimate of $\mathrm{df}(\hat{\xi})$. This result is directly applicable in practice and provides fast and accurate risk estimation without the need for cross validation, say.

Our main result is derived under the assumptions on the estimator as stated below. To fix notation we let $B(x, r)$ denote the closed ball of radius r and center x. Additionally, we let $\mathcal{H}^{n-1}$ denote the $n-1$ dimensional Hausdorff measure – a generalisation of the surface measure of $n-1$ dimensional hypersurfaces in $\mathbb{R}^n$ (see e.g. Evans & Gariepy (1992) for details).

Assumption 2.2. *The estimator $\hat{\xi}$ can be written as $\hat{\xi} = \sum_{i=1}^N 1_{U_i}\hat{\xi}_i$ for a collection of open and disjoint sets $\{U_i\}_{i=1}^N$ with $\bigcup_{i=1}^N \overline{U}_i = \mathbb{R}^n$ and $\mathrm{int}(\overline{U}_i) = U_i$. Additionally, for each $i = 1, ..., n$:*

(a) The map $\hat{\xi}_i : \overline{U}_i \to \mathbb{R}^n$ is locally Lipschitz.

(b) The random variable $1_{U_i}\mathrm{div}(\hat{\xi}_i)(Y)$ has finite first moment and $\|\hat{\xi}_i\|$ is polynomially bounded on U_i.

(c) The function $r \mapsto \mathcal{H}^{n-1}(\partial U_i \cap B(0, r))$ is polynomially bounded.

Remark 2.3. The following points are worth noting:

a) **Boundary values of the estimator.** Assumption 2.2(c) implies that the boundaries of the sets U_i are Lebesgue null sets, and thus that $\mathbb{R}^n \setminus \bigcup_i U_i$ has Lebesgue measure zero. The estimator $\hat{\xi}$ is here defined to be zero on this null set, but with Y having an absolutely continuous distribution its value on a null set is irrelevant. Note, however, that Assumption 2.2(a) ensures that $\hat{\xi}_i$ is uniquely defined on ∂U_i. In a concrete case there may be a natural way to define $\hat{\xi}$ on the common boundary between U_i and U_j, say, but we make no abstract attempt to select between ξ_i and ξ_j on the boundary.

b) **Degrees of freedom.** Assumption 2.2(a) implies by Rademacher's theorem (Theorem 3.1.6 and 3.1.7 in Federer (1969)) that $\mathrm{div}(\hat{\xi}_i)$ is defined Lebesgue a.e.. Combining this with Assumption 2.2(b) we conclude that under Assumption 2.2 both $\mathrm{df}(\hat{\xi})$ and $\mathrm{df}_{\mathrm{S}}(\hat{\xi})$ are well defined.

c) **Existence of normal vectors.** Assumption 2.2(c) implies that the sets U_i have locally finite perimeter (see Theorem 5.11.1 in Evans & Gariepy (1992)), thus a measure theoretic outer unit normal η_i is defined on a subset of ∂U_i (Lemma 5.8.1 in Evans & Gariepy (1992)). By convention, η_i is set equal to zero when it is not defined. Whenever ∂U_i is smooth the measure theoretic unit normal coincides with the usual pointwise unit normal.

d) **Regular open assumption.** The sets U_i are assumed regular open, that is, $\operatorname{int}(\overline{U}_i) = U_i$, which among other things implies that $\partial U_i \subseteq \bigcup_{j:j\neq i} \overline{U}_j$. For the details see the proof of Lemma A.2. It is a convenient though not strictly necessary assumption. Given Assumptions 2.2(a) and 2.2(c) the estimator is only changed on a Lebesgue null set if U_i is replaced by $\operatorname{int}(\overline{U}_i)$ on which $\hat{\xi}_i$ is well defined by Assumption 2.2(a). This can only make the boundaries smaller and Assumption 2.2(c) still holds after the replacement.

Estimators that involve data driven variable selection will generally fulfill Assumption 2.2 with each U_i corresponding to a set of selected variables. Example 2.5 and Appendix A.1 provide a thorough characterization of U_i in the debiased lasso setup. Moreover, a similar characterization of U_i is given in Example 3.4 for a class of estimators defined via minimization of a penalized loss function.

The conditions in Assumption 2.2 are typically easy to verify, except perhaps the third condition, as it involves bounding Hausdorff measures. Appendix A.2 provides some results that can be helpful for verifying the third condition. For estimators satisfying Assumption 2.2 we have the following representation of the degrees of freedom.

Theorem 2.4. *If $\hat{\xi}$ satisfies Assumption 2.2 then*

$$\operatorname{df}(\hat{\xi}) = \operatorname{df}_{\mathrm{S}}(\hat{\xi}) + \frac{1}{2}\sum_{i\neq j}\int_{\overline{U}_i\cap\overline{U}_j} \langle \hat{\xi}_j - \hat{\xi}_i, \eta_i\rangle \psi(\,\cdot\,;\xi,\sigma^2)\; d\mathcal{H}^{n-1}, \tag{4}$$

where η_i denotes the measure theoretic outer unit normal to ∂U_i.

The proof is in Appendix A.3. The essential part is an application of a generalized version of Gauss-Green's formula combined with a dominated convergence argument. Note that though $\overline{U}_i \cap \overline{U}_j$ is a Lebesgue null set – on which $\hat{\xi}$ is defined to be zero – $\hat{\xi}_j$ and $\hat{\xi}_i$ are uniquely defined by Assumption 2.2(a) and generally non-zero and different, cf. also Remark 2.3(a).

If $\hat{\xi}$ satisfies Assumption 2.2 and is continuous then (4) reduces to $\operatorname{df}(\hat{\xi}) = \operatorname{df}_{\mathrm{S}}(\hat{\xi})$, which is Stein's lemma for a class of locally Lipschitz continuous estimators. The boundary integrals therefore account for potential jumps of $\hat{\xi}$ across the boundary of any two adjacent regions U_i and U_j.

The following example illustrates how to verify Assumption 2.2 for the debiased lasso estimator, which is the estimator that will also be the main focus of the subsequent section.

Example 2.5 (The debiased lasso estimator)**.** Consider linear regression estimators of the form $\hat{\xi} = X\hat{\beta}$, where X is an $n \times p$-matrix and $\hat{\beta} : \mathbb{R}^n \to \mathbb{R}^p$. We will assume throughout that X has rank p and thus $p \leq n$. For any subset $A \subseteq \{1, ..., p\}$, X_A denotes the matrix whose columns are those of X indexed by A, and similarly, β_A denotes $(\beta_i)_{i\in A}$. We denote by

$$\Pi_A = X_A(X_A^T X_A)^{-1} X_A^T$$

the orthogonal projection onto $\mathrm{col}(X_A)$ – the column space of X_A. We let $\mathcal{A}$ denote the power set of $\{1, ..., p\}$ and $\mathrm{supp}(\beta) = \{i \mid \beta_i \neq 0\} \in \mathcal{A}$ denote the support of $\beta \in \mathbb{R}^p$.

The *lasso estimator* of β with tuning parameter $\lambda > 0$ is defined as

$$\hat{\beta}^\lambda(y) = \arg\min_\beta \frac{1}{2}\|y - X\beta\|_2^2 + \lambda\|\beta\|_1.$$

It follows from the first order subgradient conditions for this convex minimization problem that $\hat{\beta}^\lambda(y)$ is the lasso estimate if and only if

$$\begin{aligned} \langle X_i, y - X\hat{\beta}^\lambda(y)\rangle &= \lambda\,\mathrm{sign}(\hat{\beta}^\lambda(y)) \quad \text{for} \quad \hat{\beta}_i^\lambda(y) \neq 0 \\ |\langle X_i, y - X\hat{\beta}^\lambda(y)\rangle| &\leq \lambda \quad \text{for} \quad \hat{\beta}_i^\lambda(y) = 0, \end{aligned}$$

see e.g. Section 3 in Tibshirani & Taylor (2012) or Lemma 2.1 in Bühlmann & van de Geer (2011). The lasso estimator of ξ is $\hat{\xi}^\lambda_{\text{lasso}} = X\hat{\beta}^\lambda$, and with $\hat{S}^\lambda(y) = \mathrm{supp}(\hat{\beta}^\lambda(y))$ denoting the set of active predictors it holds that $\mathrm{df}_S(\hat{\xi}^\lambda_{\text{lasso}}) = \mathrm{df}(\hat{\xi}^\lambda_{\text{lasso}}) = E|\hat{S}^\lambda(Y)|$, see Zou et al. (2007). Note that since the columns of X are assumed linearly independent the lasso estimate $\hat{\beta}^\lambda(y)$ is unique, hence $\hat{S}^\lambda(y)$ is well defined. The map $y \mapsto \hat{\beta}^\lambda(y)$ is well defined as a single valued map, and from Theorem 1 in Zou et al. (2007) it is continuous.

The *debiased lasso estimator* is the estimator $\hat{\xi}^\lambda_{\text{d.lasso}} = \Pi_{\hat{S}^\lambda}$, which is the orthogonal projection onto the (data driven) subspace determined by the active set $\hat{S}^\lambda$ for the lasso estimator. Introducing the sets

$$U_A^\lambda := \left\{ y \in \mathbb{R}^n \;\middle|\; |\hat{\beta}_i^\lambda(y)| > 0 \text{ for } i \in A \text{ and } |\langle X_i, y - X\hat{\beta}^\lambda(y)\rangle| < \lambda \text{ for } i \notin A \right\}$$

for $A \in \mathcal{A}$ it follows from the first order conditions that if $y \in U_A^\lambda$ then $\mathrm{supp}(\hat{\beta}^\lambda(y)) = A$. It also follows that the sets U_A^λ for $A \in \mathcal{A}$ are disjoint and from the definition that they are open. From the first order subgradient conditions we get, in addition, that $\mathbb{R}^p = \bigcup_A \overline{U}_A^\lambda$. Moreover, if $y \in \partial U_A^\lambda$ then either $\hat{\beta}_i^\lambda(y) = 0$ for some $i \in A$ or $|\langle X_i, y - X\hat{\beta}^\lambda(y)\rangle| = \lambda$ for some $i \notin A$. It is then possible to find $A' \neq A$ such that $y \in \partial U_{A'}^\lambda$ as well. This shows that $\mathrm{int}(\overline{U}_A^\lambda) = U_A^\lambda$. Observe that $\hat{\xi}^\lambda_{\text{d.lasso}} = \Pi_A$ on U_A^λ, hence

$$\hat{\xi}^\lambda_{\text{d.lasso}} = \sum_{A \in \mathcal{A}} 1_{U_A^\lambda} \Pi_A \tag{5}$$

on $\bigcup_A U_A^\lambda$, and we will henceforth take (5) as the definition of $\hat{\xi}^\lambda_{\text{d.lasso}}$ on $\mathbb{R}^n$. Since Π_A is Lipschitz continuous – and in particular polynomially bounded – with $\mathrm{div}(\Pi_A) = |A|$ we conclude that $\hat{\xi}^\lambda_{\text{d.lasso}}$ is an estimator of the form required by Assumption 2.2, which fulfills Assumptions 2.2(a) and 2.2(b). Assumption 2.2(c) remains to be demonstrated.

Since $\beta_i^\lambda(y) = 0$ for $y \in U_A^\lambda$ and $i \notin A$ it follows by continuity that $\hat{\xi}^\lambda_{\text{lasso}}(y) \in \mathrm{col}(X_A)$ for $y \in \overline{U}_A^\lambda$. Hence, for any pair $A_1, A_2 \in \mathcal{A}$, it holds that $\hat{\xi}^\lambda_{\text{lasso}}(y) \in \mathrm{col}(X_{A_1}) \cap \mathrm{col}(X_{A_2}) = \mathrm{col}(X_{A_1 \cap A_2})$ whenever $y \in \overline{U}_{A_1}^\lambda \cap \overline{U}_{A_2}^\lambda$. Moreover, by defining the affine subspaces

$$L_A(s) := \{u \in \mathbb{R}^n \mid \langle X_i, u\rangle = s_i\lambda \text{ for } i \in A\}$$

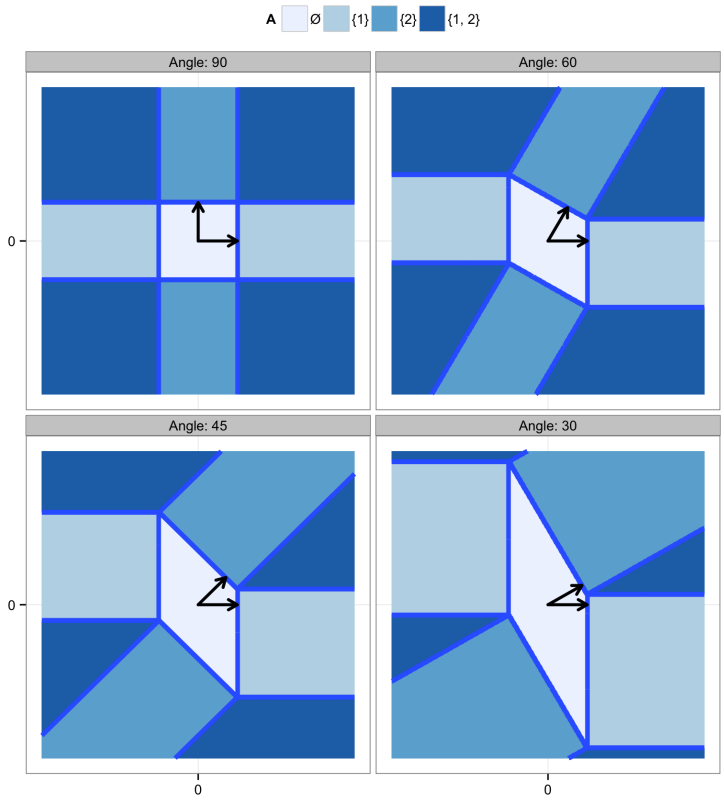

FIGURE 1. Illustrations of the decomposition of $\mathbb{R}^2$ into the four sets $U^1_\emptyset$, $U^1_{\{1\}}$, $U^1_{\{2\}}$ and $U^1_{\{1,2\}}$ according to the lasso estimator with $\lambda = 1$. The set $U^1_\emptyset$ consists of the points shrunk to zero, the sets $U^1_{\{1\}}$ and $U^1_{\{2\}}$ to the points where either the second or the first coordinate, respectively, is shrunk to zero and $U^1_{\{1,2\}}$ to the set where none of the coordinates are shrunk to zero. The decomposition depends on the angle between the two columns in X.

for $s \in \{-1, 1\}^{|A|}$, it follows from the first order conditions that

$$y - \hat{\xi}^\lambda_{\text{lasso}}(y) \in \bigcup_s L_{A_1 \cup A_2}(s)$$

for $y \in \overline{U}^\lambda_{A_1} \cap \overline{U}^\lambda_{A_2}$. Thus we have that

$$\overline{U}^\lambda_{A_1} \cap \overline{U}^\lambda_{A_2} \subseteq \operatorname{col}(X_{A_1 \cap A_2}) + \bigcup_s L_{A_1 \cup A_2}(s), \tag{6}$$

which by Remark 2.3(d) implies that

$$\partial U^\lambda_{A_1} \subseteq \bigcup_{A_2 : A_2 \neq A_1} \bigcup_s \left(\operatorname{col}(X_{A_1 \cap A_2}) + L_{A_1 \cup A_2}(s)\right). \tag{7}$$

The latter is a finite union of affine subspaces of dimensions

$$|A_1 \cap A_2| + n - |A_1 \cup A_2| = n + 2|A_1 \cap A_2| - |A_1| - |A_2| \leq n - 1.$$

Hence Assumption 2.2(c) holds. Figure 2 provides an illustration of the partition of $\mathbb{R}^n$ for $n = p = 2$ for different choices of angles between the columns in X. □

The arguments in Example 2.5 above are elementary and based on the first order conditions for the lasso estimator. It is possible to express the arguments in a more abstract but powerful way. We briefly outline these arguments in Appendix A.1.

3. Risk estimation for debiased lasso

It is not obvious how the general formula in Theorem 2.4 for $\mathrm{df}(\hat{\xi})$ can be used for computing or estimating the degrees of freedom. The first term, $\mathrm{df}_{\mathrm{S}}(\hat{\xi})$, may be estimated by $\mathrm{div}(\hat{\xi})(Y)$, but the second term is more difficult. In this section we show how this second term can be related to the derivative of $\lambda \mapsto \mathrm{df}_{\mathrm{S}}(\hat{\xi}^{\lambda}_{\mathrm{d.lasso}})$ for debiased lasso. First we recapitulate the computations in Tibshirani (2015) of the degrees of freedom for debiased lasso with X orthogonal, which will reveal the general formula shown below.

Example 3.1 (Continuation of Example 2.5)**.** Assume that $n = p$ and $X = I$. In this case it is well known that the lasso and the debiased lasso estimators become the soft and hard threshold estimators, respectively. That is,

$$\hat{\xi}^{\lambda}_{\mathrm{lasso},i} = \begin{cases} Y_i - \lambda \operatorname{sign}(Y_i) & \text{if } |Y_i| > \lambda \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad \hat{\xi}^{\lambda}_{\mathrm{d.lasso},i} = \begin{cases} Y_i & \text{if } |Y_i| > \lambda \\ 0 & \text{otherwise} \end{cases}.$$

We can write up closed form expressions for $\mathrm{df}(\hat{\xi}^{\lambda}_{\mathrm{d.lasso}})$ and $\mathrm{df}_{\mathrm{S}}(\hat{\xi}^{\lambda}_{\mathrm{d.lasso}})$:

$$\begin{aligned} \mathrm{df}_{\mathrm{S}}(\hat{\xi}^{\lambda}_{\mathrm{d.lasso}}) &= \int \psi(y; \xi, \sigma^2) \sum_i 1_{(|y_i| > \lambda)} \, dy = \sum_i \int_{(|y_i| > \lambda)} \psi(y_i; \xi_i, \sigma^2) \, dy_i \\ &= \sum_i \Phi\left(\frac{-\lambda - \xi_i}{\sigma}\right) + \left(1 - \Phi\left(\frac{\lambda - \xi_i}{\sigma}\right)\right), \end{aligned}$$

and as in Tibshirani (2015)

$$\begin{aligned} \mathrm{df}(\hat{\xi}^{\lambda}_{\mathrm{d.lasso}}) &= \sum_i \int_{\lambda}^{\infty} \psi(y_i; \xi_i, \sigma^2) \frac{y_i(y_i - \xi_i)}{\sigma^2} \, dy_i + \int_{-\infty}^{-\lambda} \psi(y_i; \xi_i, \sigma^2) \frac{y_i(y_i - \xi_i)}{\sigma^2} \, dy_i \\ &= \sum_i \left[-\psi(y_i; \xi_i, \sigma^2) y_i\right]_{\lambda}^{\infty} + \int_{\lambda}^{\infty} \psi(y_i; \xi_i, \sigma^2) \, dy_i \\ &\qquad + \left[-\psi(y_i; \xi_i, \sigma^2) y_i\right]_{-\infty}^{-\lambda} + \int_{-\infty}^{-\lambda} \psi(y_i; \xi_i, \sigma^2) \, dy_i \\ &= \lambda \sum_i \left(\psi(\lambda; \xi_i, \sigma^2) + \psi(-\lambda; \xi_i, \sigma^2)\right) + \mathrm{df}_{\mathrm{S}}(\hat{\xi}^{\lambda}_{\mathrm{d.lasso}}). \end{aligned}$$

Letting ∂_λ denote the differential operator with respect to λ we observe that

$$\mathrm{df}(\hat{\xi}^{\lambda}_{\mathrm{d.lasso}}) = \mathrm{df}_{\mathrm{S}}(\hat{\xi}^{\lambda}_{\mathrm{d.lasso}}) - \lambda \partial_\lambda \mathrm{df}_{\mathrm{S}}(\hat{\xi}^{\lambda}_{\mathrm{d.lasso}}), \tag{8}$$

which is a striking identity. This is because the formula for $\mathrm{df}(\hat{\xi}^{\lambda}_{\mathrm{d.lasso}})$, though explicit, involves the unknown parameter ξ and is not readily estimable. But we have the divergence estimator, $\sum_i 1_{(|y_i| > \lambda)}$, of $\mathrm{df}_{\mathrm{S}}(\hat{\xi}^{\lambda}_{\mathrm{d.lasso}})$, and if we from this

can estimate its derivative as well, the formula above suggests how to estimate $\mathrm{df}(\hat{\xi}^{\lambda}_{\text{d.lasso}})$. □

The remarkable fact that we will show is that (8) holds without the orthogonality assumption on X.

Theorem 3.2. *For the debiased lasso estimator defined in Example 2.5 it holds that*

$$\mathrm{df}(\hat{\xi}^{\lambda}_{\text{d.lasso}}) = \mathrm{df}_{\mathrm{S}}(\hat{\xi}^{\lambda}_{\text{d.lasso}}) - \lambda \partial_{\lambda} \mathrm{df}_{\mathrm{S}}(\hat{\xi}^{\lambda}_{\text{d.lasso}}) \tag{9}$$

where ∂_{λ} denotes differentiation w.r.t. λ.

Theorem 3.2 suggests that $\mathrm{df}(\hat{\xi}^{\lambda}_{\text{d.lasso}})$ can be estimated by differentiation of an estimate of $\mathrm{df}_{\mathrm{S}}(\hat{\xi}^{\lambda}_{\text{d.lasso}})$. The divergence estimate of Stein's degrees of freedom is, however, not differentiable as a function of λ, and we need to somehow smooth it. To this end it is convenient to reparametrize the penalization in terms of $\delta = \log(\lambda)$, so that with

$$h(\delta) \coloneqq \mathrm{df}_{\mathrm{S}}(\hat{\xi}^{\exp(\delta)}_{\text{d.lasso}}),$$

then

$$\mathrm{df}(\hat{\xi}^{\exp(\delta)}_{\text{d.lasso}}) = h(\delta) - h'(\delta).$$

In simulations h was found to be monotonely decreasing, and thus h' to be negative, but we cannot prove that this is generally the case. The integral representation of h' from Theorem 2.4 is not particularly helpful as the integrand can, in fact, be negative. Based on our computational observations – and to reduce variance of the resulting estimate – our proposal is based on the assumption that h' is negative. It is effectively a kernel smoother that estimates the intensity of jumps for a monotone jump process.

We note that $|\hat{S}^{\exp(\delta)}|$ is an unbiased estimate of $h(\delta)$ and that the function $\delta \mapsto |\hat{S}^{\exp(\delta)}|$ is a step function. The problem of estimating the derivative, h', of its mean is thus analogous to estimating the intensity for a jump process with one main difference; the step function can have jumps of negative as well as positive sign, though most jumps will be negative. Our proposed estimate ignores the positive excursions of the step function and is computed as follows:

- Compute the jump points, λ_i and jump sizes, $\Delta_i \coloneqq \inf_{\lambda < \lambda_i} |\hat{S}^{\lambda}| - |\hat{S}^{\lambda_i +}|$, of the decreasing function $\lambda \mapsto \inf_{\lambda' < \lambda} |\hat{S}^{\lambda'}|$ for $i = 1, \ldots, M$.
- Apply a kernel density smoother to the points $\delta_i = \log(\lambda_i)$ for $i = 1, \ldots, M$ counted with the multiplicities Δ_i. In the simulations presented in this paper an adaptive Gaussian kernel density smoother was used (see Section 10.4.3.2 in Givens & Hoeting (2012)).
- Rescale the density estimate by the total number of jumps, that is, by $\sum_{i=1}^{M} \Delta_i$.

As mentioned above, we can think of the proposed estimate of h' as a non-parametric estimate of the intensity of the jumps for a monotonely decreasing jump process. Alternatively, we can think of it as smoothing the jumps by a sigmoidal function (the anti-derivative of the kernel) to obtain a smooth estimate of Stein's degrees of freedom, which can then be differentiated. Note that even if Δ_i may always be 1 in theory, the jumps are in practice computed on a grid and may thus be larger than 1, which the procedure accounts for. The estimate of $-\lambda \partial_{\lambda} \mathrm{df}_{\mathrm{S}}(\hat{\xi}^{\lambda}_{\text{d.lasso}})$ resulting from the procedure above is denoted by $\widehat{\partial}$.

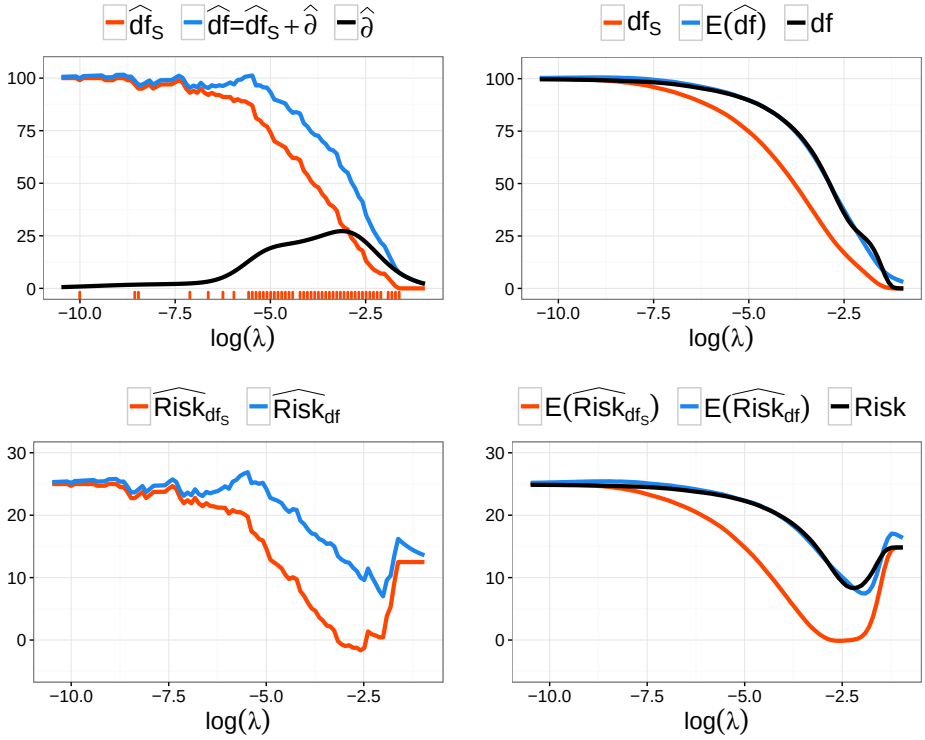

FIGURE 2. Left: Realization of the estimates of degrees of freedom $\hat{\mathrm{df}}_S = |\hat{S}^\lambda|$ and $\hat{\mathrm{df}} = |\hat{S}^\lambda| + \widehat{\partial}$ as well as the correction term $\widehat{\partial}$ as a function of $\log(\lambda)$ (top) and corresponding estimates of the risk (bottom). Right: Similar to the left but mean values of the estimates obtained by averaging over 1000 samples along with the degrees of freedom $\mathrm{df} = \mathrm{df}(\hat{\xi}^\lambda_{\mathrm{d.lasso}})$ obtained from the 1000 samples using the covariance definition (1). The design parameters were: $\sigma = 0.5$, $n = p = 100$, $\gamma = 1$, $\alpha = 0.1$ and the design type was (S) with constant correlation of $\rho = 0.1$ (see Section 4).

Using $|\hat{S}^\lambda| + \widehat{\partial}$ as an estimate of degrees of freedom leads to the risk estimate

$$\widehat{\mathrm{Risk}}_{\mathrm{df}} := \|Y - \hat{\xi}^\lambda_{\mathrm{d.lasso}}\|_2^2 - n\sigma^2 + 2\sigma^2 \left(|\hat{S}^\lambda| + \widehat{\partial}\right). \tag{10}$$

For an example of the above estimate see Figure 2, where $\widehat{\partial}$ and $\widehat{\mathrm{Risk}}_{\mathrm{df}}$ are applied to a single realization of Y along with an average over 1000 replications.

To prove Theorem 3.2 we prove a more general intermediate result for estimators that are parametrized in a similar way by a tuning parameter. We use in the following D to denote the differential operator w.r.t. y.

Proposition 3.3. *Let $q > 0$ and suppose that $\hat{\xi}^\lambda = \sum_i 1_{U_i^\lambda} \hat{\xi}_i$ where*

$$U_i^\lambda = \lambda^q U_i^1, \qquad \textit{for all } i = 1, ..., N. \tag{11}$$

Assume that $\mathrm{div}(\hat{\xi}_i)$ *is locally Lipschitz and both* $\mathrm{div}(\hat{\xi}_i)$ *and* $D(\mathrm{div}(\hat{\xi}_i))$ *are polynomially bounded for each* $i = 1, ..., N$. *If* $\hat{\xi}^1$ *satisfies Assumption 2.2 then*

(12)
$$-\frac{\lambda}{q}\partial_\lambda \mathrm{df}_S(\hat{\xi}^\lambda) = \frac{1}{2}\sum_{i\neq j}\int_{\overline{U}_i^\lambda \cap \overline{U}_j^\lambda} \Big(\mathrm{div}(\hat{\xi}_j)(y) - \mathrm{div}(\hat{\xi}_i)(y)\Big)\langle y, \eta_i\rangle \psi(y;\xi,\sigma^2)\; d\mathcal{H}^{n-1}(y).$$

Proof. First observe that $\partial U_i^\lambda \cap B(0,r) = \partial U_i^1 \cap B(0, r/\lambda^q)$, hence if $\hat{\xi}^1$ satisfies Assumption 2.2 so does $\hat{\xi}^\lambda$ for all λ. Next, the change of variable formula yields

$$\begin{aligned}\mathrm{df}_S(\hat{\xi}^\lambda) &= \int \psi(y)\mathrm{div}(\hat{\xi}^\lambda)(y)\, dy = \sum_i \int_{U_i^\lambda} \psi(y)\mathrm{div}(\hat{\xi}_i)(y)\, dy \\ &= \sum_i \int_{U_i^1} \lambda^{qn}\left(\psi \mathrm{div}(\hat{\xi}_i)\right)(\lambda^q z)\; dz.\end{aligned}$$

Here $\psi = \psi(\cdot;\xi,\sigma^2)$ to ease notation.

The last integrand is differentiable w.r.t. λ (for Lebesgue a.a. z) and its derivative is

$$\begin{aligned}&qn\lambda^{qn-1}\left(\psi\mathrm{div}(\hat{\xi}_i)\right)(\lambda^q z) + \lambda^{qn}\left\langle D\left(\psi\mathrm{div}(\hat{\xi}_i)\right)(\lambda^q z), q\lambda^{q-1}z\right\rangle \\ &\quad = \frac{q}{\lambda}\lambda^{qn}\left(n\left(\psi\mathrm{div}(\hat{\xi}_i)\right)(\lambda^q z) + \left\langle D\left(\psi\mathrm{div}(\hat{\xi}_i)\right)(\lambda^q z), \lambda^q z\right\rangle\right),\end{aligned}$$

which is integrable due to the polynomial bounds. Hence, by the change of variable formula

$$\begin{aligned}\frac{\lambda}{q}\partial_\lambda \mathrm{df}_S(\hat{\xi}^\lambda) &= \sum_i \int_{U_i^1} \lambda^{qn}\left(n\left(\psi\mathrm{div}(\hat{\xi}_i)\right)(\lambda^q z) + \left\langle D\left(\psi\mathrm{div}(\hat{\xi}_i)\right)(\lambda^q z), \lambda^q z\right\rangle\right) dz \\ &= \sum_i \int_{U_i^\lambda} n\left(\psi\mathrm{div}(\hat{\xi}_i)\right)(y) + \left\langle D\left(\psi\mathrm{div}(\hat{\xi}_i)\right)(y), y\right\rangle dy \\ &= \sum_i \int_{U_i^\lambda} n\left(\psi\mathrm{div}(\hat{\xi}_i)\right)(y) + \left\langle \left(\psi D\mathrm{div}(\hat{\xi}_i) + \mathrm{div}(\hat{\xi}_i)D\psi\right)(y), y\right\rangle\; dy \\ &= \sum_i \int_{U_i^\lambda} \psi(y)\mathrm{div}\left(y\mathrm{div}(\hat{\xi}_i)(y)\right) + \left\langle D\psi(y), y\mathrm{div}(\hat{\xi}_i)(y)\right\rangle dy.\end{aligned}$$

The last line is identified as $\mathrm{df}_S(\tilde{\xi}^\lambda) - \mathrm{df}(\tilde{\xi}^\lambda)$, where

$$\tilde{\xi}^\lambda(y) := \sum_i 1_{U_i^\lambda}(y) y \mathrm{div}(\hat{\xi}_i)(y).$$

Finally (12) follows by applying Theorem 2.4 to $\tilde{\xi}^\lambda$ (which also satisfies Assumption 2.2). □

Example 3.4. There are naturally occurring examples besides the lasso generated sets that satisfy (11). Consider still a linear regression setup with X an $n \times p$-matrix of rank p. Let ℓ denote the penalized loss function

$$\ell(y,\beta,\lambda) = \frac{1}{2}\|y - X\beta\|_2^2 + \lambda \mathrm{Pen}(\beta),$$

for some penalty function $\mathrm{Pen} : \mathbb{R}^p \to \mathbb{R}$ and define

$$U_A^\lambda = \mathrm{int}\left\{ y \in \mathbb{R}^n \;\middle|\; \inf_{\beta:\mathrm{supp}(\beta)=A} \ell(y,\beta,\lambda) = \inf_\beta \ell(y,\beta,\lambda)\right\}, \tag{13}$$

for each $A \in \mathcal{A}$. Hence any $y \in U_A^\lambda$ has X_A as its active predictors. If Pen is *positive homogeneous* of degree $k \in [0, 2)$ then

$$\ell\left(\lambda^{\frac{1}{2-k}} y, \lambda^{\frac{1}{2-k}} \beta, \lambda\right) = \lambda^{\frac{2}{2-k}} \ell\left(y, \beta, 1\right).$$

Hence $U_A^\lambda = \lambda^{\frac{1}{2-k}} U_A^1$ holds for all $A \in \mathcal{A}$ and $\lambda > 0$. The k-norms, $\mathrm{Pen}(\beta) = \|\beta\|_k^k$ for $k \in (0, 2)$ and $\mathrm{Pen}(\beta) = \|\beta\|_0 = |\mathrm{supp}(\beta)|$ for $k = 0$, are examples of positive homogeneous penalties. For these penalties only $k \in [0, 1]$ will result in variable selection. With $\mathrm{Pen}(\cdot) = \|\cdot\|_1$ we see that for lasso the sets U_A^λ satisfy (11) with $q = 1$.

□

Proof of Theorem 3.2. Let U_A^λ be defined as in Example 2.5, where it was also shown that Assumption 2.2 holds for the debiased lasso estimator. Moreover, from Example 3.4 we see that $U_A^\lambda = \lambda U_A^1$ for all $\lambda > 0$ and $A \in \mathcal{A}$. By Theorem 2.4 we know that the left hand side of (9) is

$$\begin{aligned} &\mathrm{df}(\hat{\xi}^\lambda_{\text{d.lasso}}) - \mathrm{df}_{\mathrm{S}}(\hat{\xi}^\lambda_{\text{d.lasso}}) \\ &\qquad = \frac{1}{2} \sum_{A_1 \neq A_2} \int_{\overline{U}^\lambda_{A_1} \cap \overline{U}^\lambda_{A_2}} \langle (\Pi_{A_2} - \Pi_{A_1}) y, \eta_{A_1}(y) \rangle \psi(y) \; d\mathcal{H}^{n-1}(y). \end{aligned} \tag{14}$$

Due to (6) and a dimensionality argument the Hausdorff measure of $\overline{U}^\lambda_{A_1} \cap \overline{U}^\lambda_{A_2}$ is zero unless $A_1 \subseteq A_2$ and $|A_2| = |A_1| + 1$ or vice versa. Furthermore, one can deduce from (6) that if $A_1 \subseteq A_2$ then η_{A_1} must belong to $\mathrm{col}(X_{A_2}) \ominus \mathrm{col}(X_{A_1})$, which is 1-dimensional when $|A_2| = |A_1| + 1$. Especially, we have $\eta_{A_1} = (\Pi_{A_2} - \Pi_{A_1}) \eta_{A_1}$, hence (14) becomes

$$\begin{aligned} \mathrm{df}(\hat{\xi}^\lambda_{\text{d.lasso}}) - \mathrm{df}_{\mathrm{S}}(\hat{\xi}^\lambda_{\text{d.lasso}}) &= \sum_{\substack{A_1 \subseteq A_2, \\ |A_2| = |A_1| + 1}} \int_{\overline{U}^\lambda_{A_1} \cap \overline{U}^\lambda_{A_2}} \langle y, \eta_{A_1}(y) \rangle \psi(y) \; d\mathcal{H}^{n-1}(y) \\ &= \sum_{\substack{A_1 \subseteq A_2, \\ |A_2| = |A_1| + 1}} \int_{\overline{U}^\lambda_{A_1} \cap \overline{U}^\lambda_{A_2}} \underbrace{\mathrm{div}(\Pi_{A_2} - \Pi_{A_1})}_{= |A_2| - |A_1| = 1} \langle y, \eta_{A_1}(y) \rangle \psi(y) \; d\mathcal{H}^{n-1}(y) \\ &= -\lambda \partial_\lambda \mathrm{df}_{\mathrm{S}}(\hat{\xi}^\lambda_{\text{d.lasso}}) \end{aligned}$$

by Proposition 3.3. □

4. Simulation Study

The purpose of the simulation study reported in this section was to quantify how the risk estimator $\widehat{\mathrm{Risk}}_{\mathrm{df}}$ given by (10) performs when used for selection of the penalty parameter λ. The risk of the resulting estimator was computed via simulations for a range of designs and parameters. Afterwards it was compared to the risk of the oracle estimator as well as the risk of the estimator obtained by selecting λ using

$$\widehat{\mathrm{Risk}}_{\mathrm{df_S}} = \|Y - \hat{\xi}^\lambda_{\text{d.lasso}}\|_2^2 - n\sigma^2 + 2\sigma^2 |\hat{S}^\lambda|, \tag{15}$$

which does not adjust for the variable selection performed by debiased lasso. The variance parameter σ^2 was considered known as well as unknown, and in the latter case it was thus estimated according to the method outlined below.

Letting $\hat{\lambda}_{\mathrm{df}}$ denote the value of λ that minimizes $\widehat{\mathrm{Risk}}_{\mathrm{df}}$, the risk of the resulting estimator is denoted

$$\mathrm{Risk}(\lambda_{\mathrm{df}}) \coloneqq E\|\xi - \hat{\xi}^{\hat{\lambda}_{\mathrm{df}}}_{\mathrm{d.lasso}}\|_2^2.$$

Likewise, $\mathrm{Risk}(\lambda_{\mathrm{df_S}})$ denotes the risk when λ is selected by minimizing $\widehat{\mathrm{Risk}}_{\mathrm{df_S}}$. If the true mean is $\xi = X_A\beta$ for $\beta \in \mathbb{R}^A$ the oracle estimator is Π_A, whose risk is

$$E\|\xi - \Pi_A Y\|_2^2 = \sigma^2|A|.$$

In the simulation study we computed $\mathrm{Risk}(\lambda_{\mathrm{df}})/(\sigma^2 n)$ and $\mathrm{Risk}(\lambda_{\mathrm{df_S}})/(\sigma^2 n)$, which can then be compared to $|A|/n$ – the fraction of nonzero parameters.

In the simulation study the mean was given as $X\beta$ with

$$\beta_i = \begin{cases} \gamma^{i-1} & \text{if } i \le \lceil n\alpha \rceil \\ 0 & \text{otherwise} \end{cases}$$

for different choices of the dimension n, the $n \times n$ design matrix X and the parameters γ and α. The following choices of parameters were considered:

parameter	values				
σ	0.1	0.2	0.5	1	2
α	0	0.05	0.1	0.3	0.5
n	100	200			
γ	1	0.9			
X	O	S	E		
ρ	0	0.1	0.4	0.7	

The parameter ρ and the values of the design require some explanation. The three different design types are:

- Orthogonal (O), where $X = I$.
- Simulated (S), where the columns of X are standard normally distributed with one of the following correlation structures:
  - Autoregressive setup: $\mathrm{corr}(X_i, X_j) = \rho^{|i-j|}$ for all $i \neq j$.
  - Constant correlation setup: $\mathrm{corr}(X_i, X_j) = \rho$ for all $i \neq j$.
- Empirical (E), where the rows and columns are randomly selected from the 240×377 matrix of microRNA expression values as used in the earlier study by Vincent et al. (2014).

The columns of the simulated and empirical designs were standardized to have norm one to obtain a comparable signal-to-noise ratio across the three designs.

The risk estimates were based on 1000 samples for each combination of the parameters, which were generated as follows. For each of the 1000 samples a design matrix X was created/simulated and a single realization of $Y \sim \mathcal{N}(X\beta, \sigma^2 I_n)$ was drawn. For each sample the losses $\|\xi - \hat{\xi}^{\hat{\lambda}_{\mathrm{df}}}_{\mathrm{d.lasso}}\|_2^2$ and $\|\xi - \hat{\xi}^{\hat{\lambda}_{\mathrm{df_S}}}_{\mathrm{d.lasso}}\|_2^2$ were computed, and the risks were estimated as the average of the losses over the 1000 samples.

All simulations were carried out assuming either that σ^2 was known or using the following estimator of σ^2: first the lasso path $\lambda \mapsto \hat{\xi}_{\mathrm{lasso}}(\lambda)$ was calculated, then $\hat{\lambda}$ was selected by minimizing the generalized cross validation criterion

$$\mathrm{gcv}(\lambda) = \frac{\|Y - \hat{\xi}_{\mathrm{lasso}}(\lambda)\|_2^2}{\left(1 - \frac{|S^\lambda|}{n}\right)^2},$$

and σ^2 was finally estimated as

$$\hat{\sigma}^2 = \frac{\|Y - \hat{\xi}^{\hat{\lambda}}_{\text{lasso}}\|_2^2}{n - |S^{\hat{\lambda}}|}.$$

The main reason for choosing this estimator was computational efficiency, as the lasso path must be calculated for debiased lasso anyway. Thus this variance estimate has virtually no extra computational costs. See also Reid et al. (2015) for a comprehensive comparison of variance estimators. Throughout, the R package *glmnet*, Friedman et al. (2010), was used to compute the lasso path.

Figure 3 shows the simulation results for $\gamma = 0.9$ and $n = 200$ for the orthogonal and empirical designs. The supplementary material contains the results for all the simulations. Some of the results are summarized in Figures 4 and 5. Tuning λ by using $|\hat{S}^{\lambda}| + \widehat{\partial}$ as an estimate of degrees of freedom is generally superior to using $|\hat{S}^{\lambda}|$ and in the worst cases at least comparable. The differences are biggest for the lowest signal-to-noise ratios. The benefit of using $|\hat{S}^{\lambda}| + \widehat{\partial}$ generally increases with the dimension n, and it increases with decreasing signal-to-noise ratio. Furthermore, when the number of non-zero parameters is large and the signal-to-noise ratio is low (specifically, $\gamma = 0.9$, α large and σ large), $\hat{\xi}^{\hat{\lambda}_{\text{df}}}_{\text{d.lasso}}$ clearly outperforms the oracle estimator, while $\hat{\xi}^{\hat{\lambda}_{\text{df}_S}}_{\text{d.lasso}}$ is comparable or worse than the oracle estimator. Neither of the estimators performs well for small variances and large signal-to-noise ratios. For the orthogonal design the estimation of the variance incurs a clear performance loss, which is not the case for the other designs. We ascribe this to the variance estimator being particularly poor for the orthogonal design.

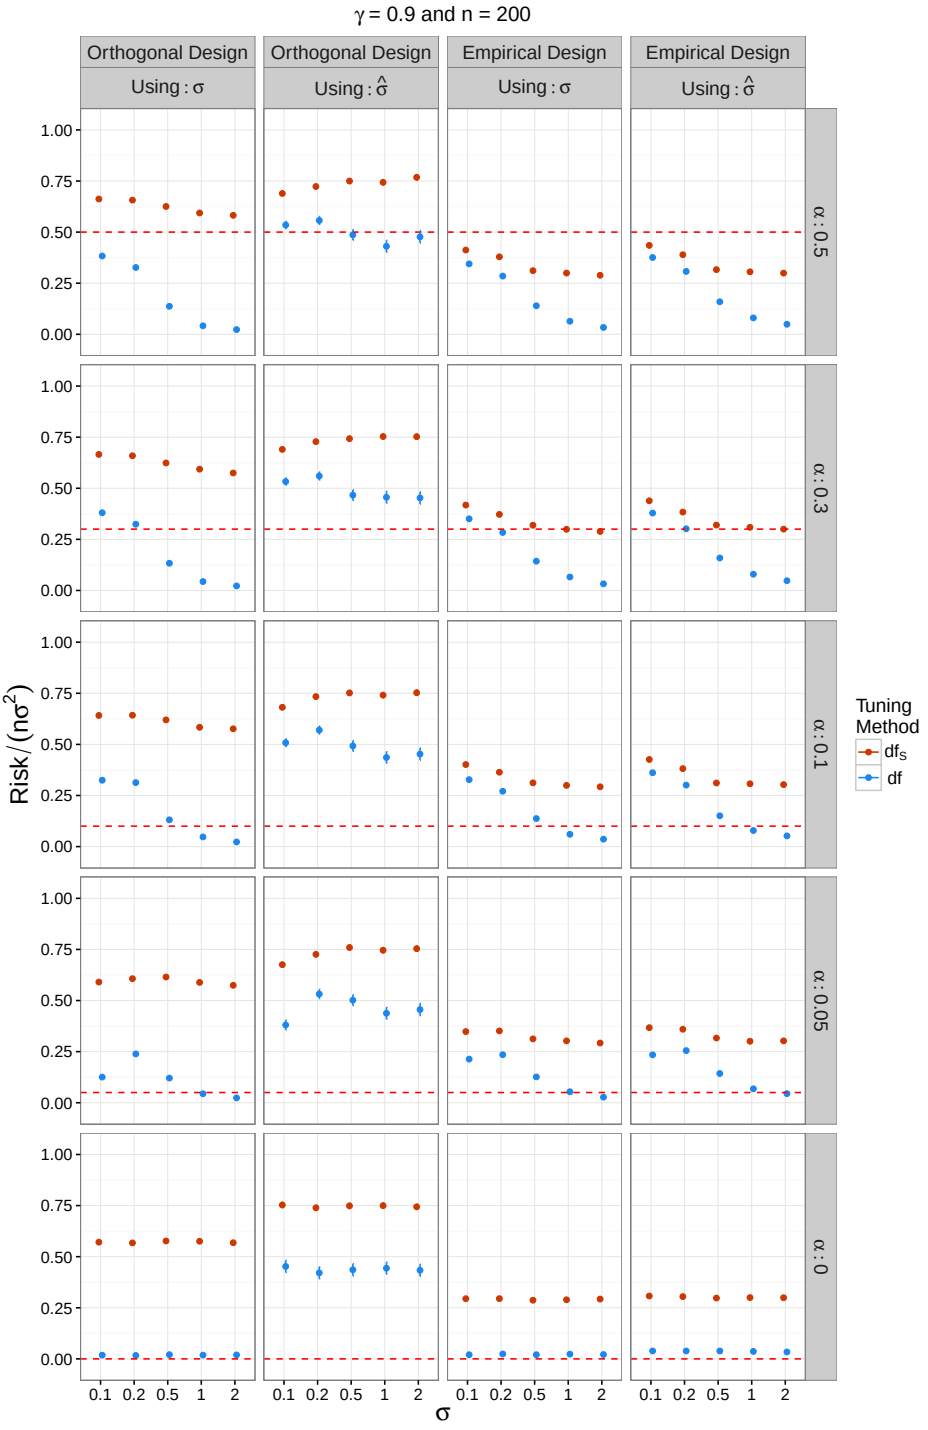

FIGURE 3. Errorbar plot of the risk estimates relative to $\sigma^2 n$ for the estimators $\hat{\xi}^{\hat{\lambda}_{\mathrm{df_S}}}_{\mathrm{d.lasso}}$ and $\hat{\xi}^{\hat{\lambda}_{\mathrm{df}}}_{\mathrm{d.lasso}}$. The dashed line is $\lceil n\alpha \rceil / n \simeq \alpha$, the relative risk for the oracle estimator.

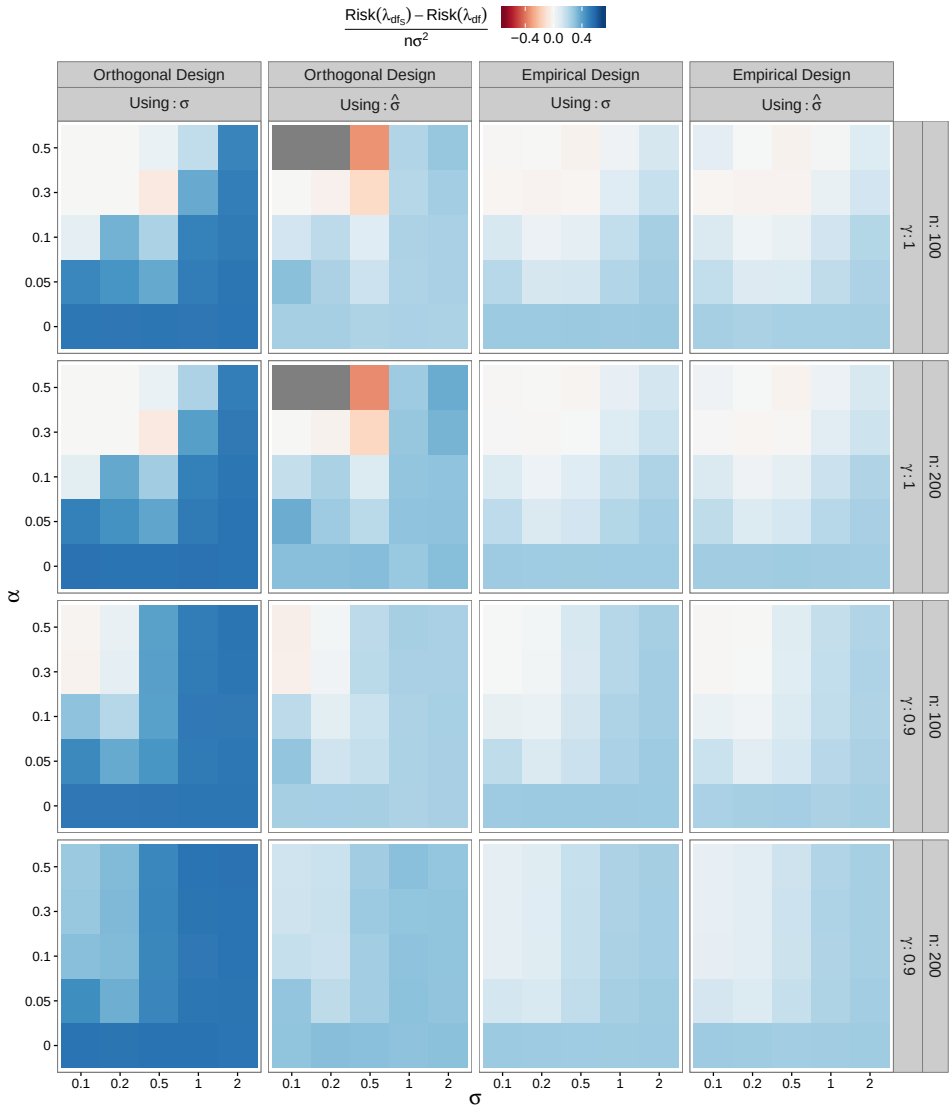

FIGURE 4. Heatmap of the difference between the average risk estimates for the estimators $\hat{\xi}^{\hat{\lambda}_{\mathrm{df_S}}}_{\mathrm{d.lasso}}$ and $\hat{\xi}^{\hat{\lambda}_{\mathrm{df}}}_{\mathrm{d.lasso}}$ for the orthogonal and empirical designs. A few configurations yielded large risk estimates (both methods exceeding $2n\sigma^2$) and large differences, these are coloured gray. The numerical values of the differences are tabulated in the supplementary material.

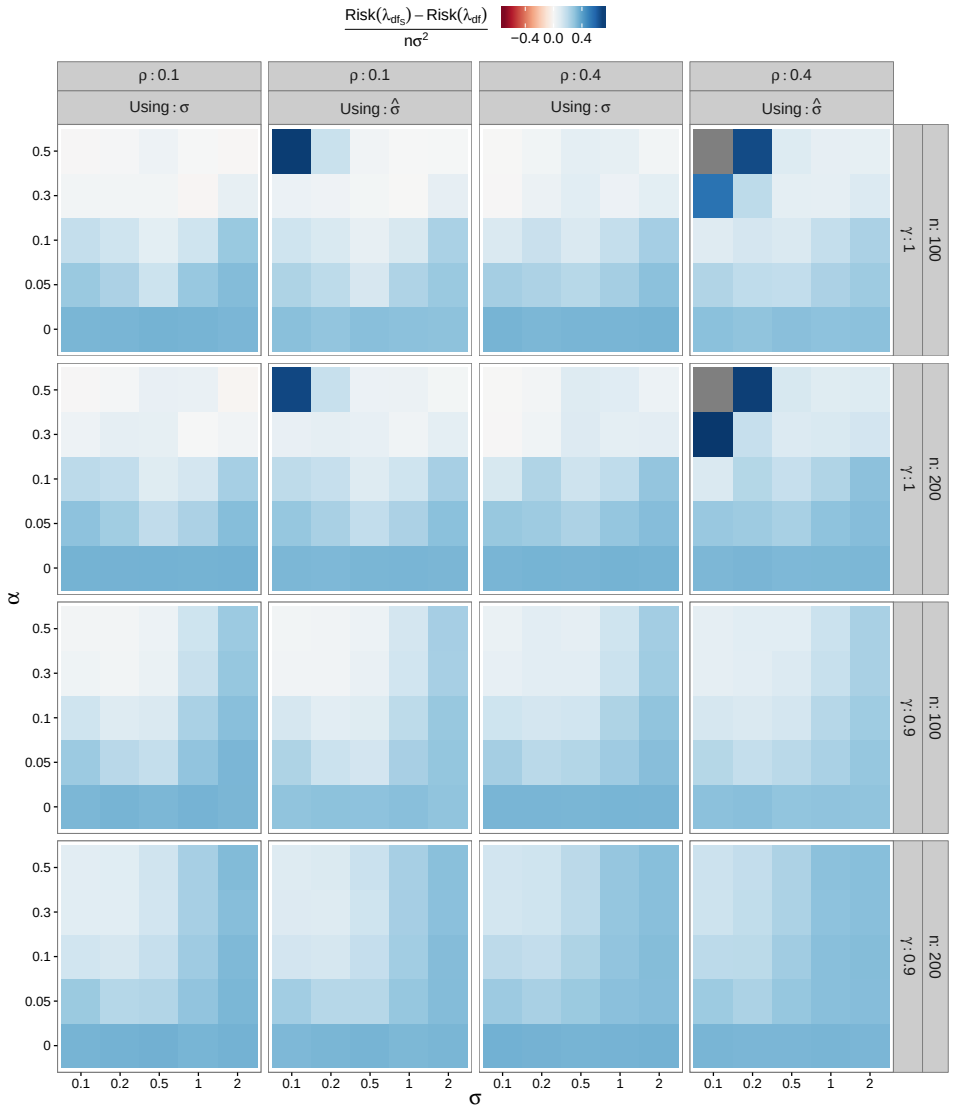

FIGURE 5. Heatmap of the differences between the average risk estimates for the estimators $\hat{\xi}^{\hat{\lambda}_{\mathrm{df_S}}}_{\mathrm{d.lasso}}$ and $\hat{\xi}^{\hat{\lambda}_{\mathrm{df}}}_{\mathrm{d.lasso}}$ for the simulated design with fixed correlation and $\rho \in \{0.1, 0.4\}$. The one configuration for which risk estimates were large (exceeding $2n\sigma^2$ for both methods) is coloured gray. The numerical values of the differences are tabulated in the supplementary material.

5. Best Subset Selection

Example 3.4 demonstrates that (11) holds for other estimators than debiased lasso, and Theorem 3.3 holds, in particular, for best subset selection in the Lagrangian formulation, which corresponds to $\mathrm{Pen}(\cdot) = \|\cdot\|_0$ being the 0-norm in Example 3.4. Theorem 3.2 does, however, only partly extend to best subset selection. In this section we demonstrate that this may still provide a practically useful estimate of degrees of freedom.

The best subset selection estimator of ξ with tuning parameter $\lambda > 0$, denoted by $\hat{\xi}^\lambda_{\mathrm{bs}}$, is

$$\hat{\xi}^\lambda_{\mathrm{bs}} = X\hat{\beta}^\lambda \quad \text{where} \quad \hat{\beta}^\lambda = \arg\min_\beta \frac{1}{2}\|Y - X\beta\|_2^2 + \lambda\|\beta\|_0.$$

It can be written on the form $\hat{\xi}^\lambda_{\mathrm{bs}} = \sum_{A\in\mathcal{A}} 1_{U^\lambda_A}\Pi_A$ (Lebesgue a.e.), where

$$U^\lambda_A := \left\{ y \in \mathbb{R}^n \;\middle|\; \lambda|A| - \frac{1}{2}\|\Pi_A y\|_2^2 < \min_{B\in\mathcal{A}\setminus A} \lambda|B| - \frac{1}{2}\|\Pi_B y\|_2^2 \right\}, \qquad A \in \mathcal{A}. \tag{16}$$

It is straight forward to verify that $\hat{\xi}^\lambda_{\mathrm{bs}}$ fulfills Assumption 2.2 except 2.2(c), which follows by Lemma A.1 in the appendix. Hence Theorem 2.4 applies to $\hat{\xi}^\lambda_{\mathrm{bs}}$.

From (16) we note that the outer unit normal to $\partial U^\lambda_{A_1}$ on $\overline{U}^\lambda_{A_1} \cap \overline{U}^\lambda_{A_2}$ equals $(\Pi_{A_2} - \Pi_{A_1})y$ normalized to have norm 1. Theorem 2.4 yields

$$\begin{aligned}
\mathrm{df}(\hat{\xi}^\lambda_{\mathrm{bs}}) - \mathrm{df}_{\mathrm{S}}(\hat{\xi}^\lambda_{\mathrm{bs}}) &= \frac{1}{2}\sum_{A_1\neq A_2} \int_{\overline{U}^\lambda_{A_1}\cap\overline{U}^\lambda_{A_2}} \frac{\langle (\Pi_{A_2} - \Pi_{A_1})y, (\Pi_{A_2} - \Pi_{A_1})y\rangle}{\|(\Pi_{A_2} - \Pi_{A_1})y\|_2}\psi(y)\; d\mathcal{H}^{n-1}(y) \\
&= \frac{1}{2}\sum_{A_1\neq A_2} \int_{\overline{U}^\lambda_{A_1}\cap\overline{U}^\lambda_{A_2}} \|(\Pi_{A_2} - \Pi_{A_1})y\|_2\,\psi(y)\; d\mathcal{H}^{n-1}(y),
\end{aligned}$$

while Proposition 3.3 and Example 3.4 yields

$$-2\lambda\partial_\lambda \mathrm{df}_{\mathrm{S}}(\hat{\xi}^\lambda_{\mathrm{bs}}) = \frac{1}{2}\sum_{A_1\neq A_2} \int_{\overline{U}^\lambda_{A_1}\cap\overline{U}^\lambda_{A_2}} \psi(y)\frac{y\cdot(\Pi_{A_2} - \Pi_{A_1})y}{\|(\Pi_{A_2} - \Pi_{A_1})y\|_2}(|A_2| - |A_1|)\; d\mathcal{H}^{n-1}(y).$$

For $A_1 \subseteq A_2$ and $|A_2| = |A_1| + 1$, we see that the integrands in the two identities above coincide. Hence, if we define

$$\mathcal{A}_1 := \left\{ (A_1, A_2) \in \mathcal{A}^2 \;\middle|\; A_1 \subseteq A_2 \text{ and } |A_2| = |A_1| + 1 \right\} \quad \text{and}$$

$$\mathcal{A}_2 := \left\{ (A_1, A_2) \in \mathcal{A}^2 \;\middle|\; A_1 \neq A_2 \text{ and } \begin{array}{c}(A_1, A_2) \notin \mathcal{A}_1 \\ (A_2, A_1) \notin \mathcal{A}_1\end{array} \right\},$$

then

$$\mathrm{df}(\hat{\xi}^\lambda_{\mathrm{bs}}) - \mathrm{df}_{\mathrm{S}}(\hat{\xi}^\lambda_{\mathrm{bs}}) = -2\lambda\partial_\lambda \mathrm{df}_{\mathrm{S}}(\hat{\xi}^\lambda_{\mathrm{bs}}) + R$$

where

$$R = \frac{1}{2}\sum_{(A_1,A_2)\in\mathcal{A}_2} \int_{\overline{U}^\lambda_{A_1}\cap\overline{U}^\lambda_{A_2}} \frac{\langle (\Pi_{A_2} - \Pi_{A_1})y, \left(\Pi_{A_2} - \Pi_{A_1} - (|A_2| - |A_1|)I_n\right) y\rangle}{\|(\Pi_{A_2} - \Pi_{A_1})y\|_2}\psi(y)\; d\mathcal{H}^{n-1}(y)$$

The usefulness of this hinges on R being small. For X orthogonal we have already demonstrated that $R = 0$ as $\hat{\xi}^\lambda_{\mathrm{bs}}$ then coincides with debiased lasso, and in this case $\overline{U}^\lambda_{A_1} \cap \overline{U}^\lambda_{A_2}$ has Hausdorff measure zero for all $(A_1, A_2) \in \mathcal{A}_2$. For non-orthogonal X this is no longer true, see Figure 7. For best subset selection there will generally be boundaries of non-zero Hausdorff measure between many more of the sets $\overline{U}^\lambda_A$

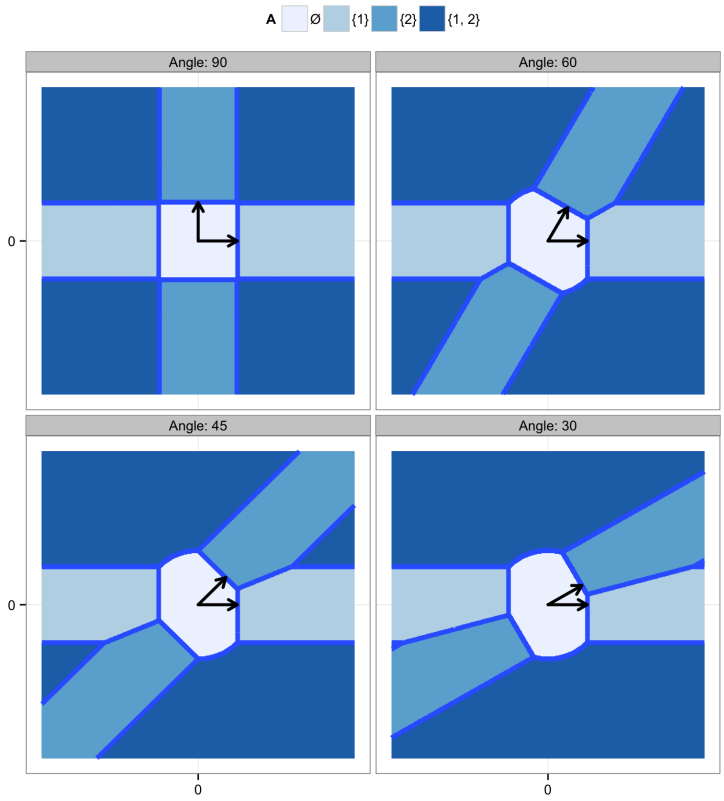

FIGURE 6. Illustrations of the decomposition of $\mathbb{R}^2$ into the four sets $U^1_\emptyset$, $U^1_{\{1\}}$, $U^1_{\{2\}}$ and $U^1_{\{1,2\}}$ according to the best subset selection estimator in the Lagrangian formulation with $\lambda = 1$. The set $U^1_\emptyset$ consists of the points projected onto the 0-dimensional space $\{0\}$, the sets $U^1_{\{1\}}$, $U^1_{\{2\}}$ to the projections onto one of the two 1-dimensional subspaces and $U^1_{\{1,2\}}$ to the identity map. The decomposition depends on the angle between the two columns in X.

– boundaries that correspond to including or excluding more than one predictor at the time or replacing predictors. Compare this with debiased lasso and Figure 2. However, by continuity in X we have $R \to 0$ for X tending to an orthogonal matrix, and we can expect R to be small for matrices that are not too far from orthogonal matrices. Thus we expect

$$\mathrm{df}_{\mathrm{S}}(\hat{\xi}^\lambda_{\mathrm{bs}}) - 2\lambda\partial_\lambda \mathrm{df}_{\mathrm{S}}(\hat{\xi}^\lambda_{\mathrm{bs}}) \tag{17}$$

to be a useful approximation for $\mathrm{df}(\hat{\xi}^\lambda_{\mathrm{bs}})$ also for non-orthogonal X.

Using the same procedure for estimating the correction $-2\lambda\partial_\lambda \mathrm{df}_{\mathrm{S}}(\hat{\xi}^\lambda_{\mathrm{bs}})$ as outlined in Section 3 – using $2\widehat{\partial}$ instead of $\widehat{\partial}$ – we used simulations to investigate if (17) was actually a good approximation of $\mathrm{df}(\hat{\xi}^\lambda_{\mathrm{bs}})$. Figure 7 shows the results using the same configurations as in Figure 2, except that n was lowered to 30 due to

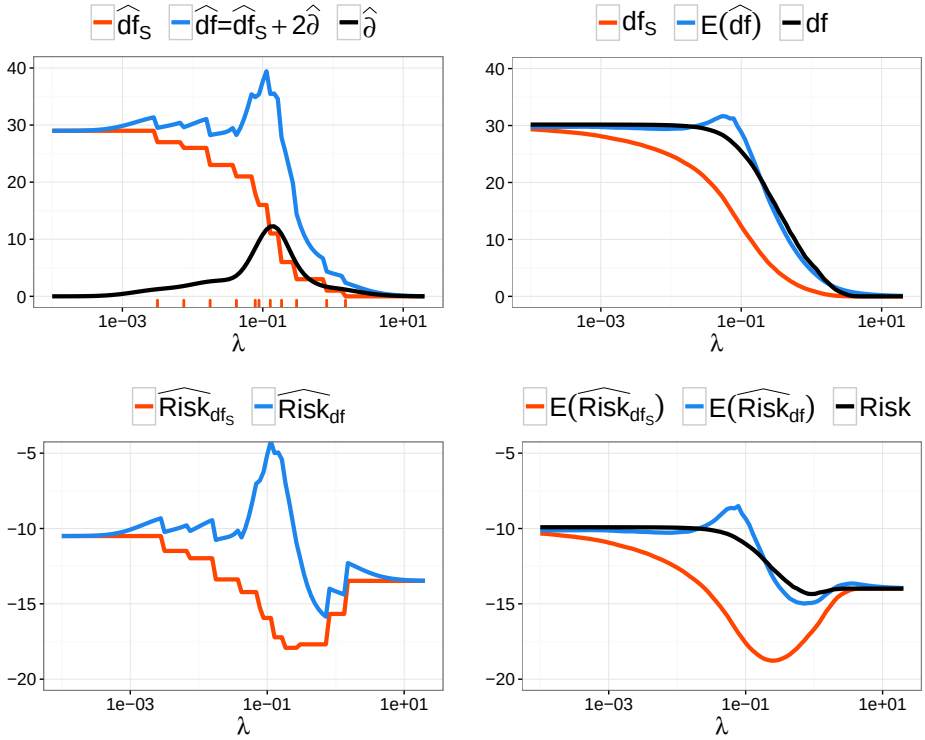

FIGURE 7. Left: Realization of the estimates of degrees of freedom $\hat{\mathrm{df}}_S = |\hat{S}^\lambda|$ and $\hat{\mathrm{df}} = |\hat{S}^\lambda| + 2\widehat{\partial}$ as well as the correction term $\widehat{\partial}$ as a function of $\log(\lambda)$ for best subset selection (top) and corresponding estimates of the risk (bottom). Right: Similar to the left but mean values of the estimates obtained by averaging over 1000 samples along with the degrees of freedom $\mathrm{df} = \mathrm{df}(\hat{\xi}^\lambda_{\mathrm{bs}})$ obtained from the 1000 samples using the covariance definition (1). The design parameters were: $\sigma = 0.5$, $n = p = 30$, $\gamma = 1$, $\alpha = 0.1$ and the design type was (S) with constant correlation of $\rho = 0.1$ (see Section 4).

computational constraints. The conclusion from this and other similar simulations (not shown) is that even with non-orthogonal designs, (17) is a practically useful approximation. That is, $-2\lambda\partial_\lambda \mathrm{df}_S(\hat{\xi}^\lambda_{\mathrm{bs}})$ accounts for the majority of the increase in the degrees of freedom due to variable selection.

6. DISCUSSION

We have provided a new representation of degrees of freedom for a broad class of discontinuous, piecewise Lipschitz estimators. This representation provides us with a deeper insight into the effect of variable selection, among other things, on the effective dimension of the statistical model and the estimator used. We have demonstrated that for debiased lasso it was, moreover, possible to derive a practically useful estimator of the degrees of freedom based on the general representation, and we also suggest that a similar estimator can be useful for best subset

selection. The estimator was based on relating the derivative of $\lambda \mapsto \mathrm{df}_S(\hat{\xi}^\lambda)$ to the discontinuities of the estimator $\hat{\xi}^\lambda$ as expressed via the integral representation of $\mathrm{df}(\hat{\xi}^\lambda) - \mathrm{df}_S(\hat{\xi}^\lambda)$. This does, indeed, make some intuitive sense as the first expresses the mean jump of degrees of freedom per unit change of λ and the other (in some sense) the mean discontinuity of degrees of freedom per unit change of y. Changing λ for fixed y or changing y for fixed λ are dual operations, and it is not surprising that we can relate the numbers.

If we were to generalize our results to other estimators that include a tuning parameter, we expect that it is only the derivative of the part of $\mathrm{df}_S(\hat{\xi}^\lambda)$ that corresponds to jumps that can be related to $\mathrm{df}(\hat{\xi}^\lambda) - \mathrm{df}_S(\hat{\xi}^\lambda)$. That is, in general, $\lambda \mapsto \mathrm{div}(\hat{\xi}^\lambda)$ will have jumps as well as smooth but non-constant pieces, and it is only the expectation of the jump part that we expect can be related to $\mathrm{df}(\hat{\xi}^\lambda) - \mathrm{df}_S(\hat{\xi}^\lambda)$. We believe that our suggested estimator of degrees of freedom may actually be generalizable to a number of discontinuous estimators involving variable selection as well as shrinkage. The requirement will be that the estimator has one or more tuning parameters and that it is computed on a grid or along a path of these. Then we can potentially estimate the derivative of the divergence of the estimator as a function of the tuning parameter(s). It is an ongoing research project to investigate this in detail.

For best subset selection we did not provide any bounds on the residual R in the approximation of $\mathrm{df}(\hat{\xi}^\lambda) - \mathrm{df}_S(\hat{\xi}^\lambda)$. It would, indeed, be very interesting to investigate this approximation in more detail. It would, in particular, be interesting to understand if it in any way can be seen as a "first order approximation" and whether there are higher order terms worth including in some cases.

Finally, we have restricted attention to the situation $n \geq p$ throughout. In many applications involving variable selection we have $n \ll p$. The general representation does hold for $n < p$ and our suggested estimator of degrees of freedom should also work for debiased lasso, but as X then no longer has linearly independent columns, some of the arguments, as those in Example 2.5, become more cumbersome. Hence for the sake of presentation we assumed $n \geq p$.

Appendix A. Additional results and Proofs

A.1. **The lasso decomposition.** The decomposition of $\mathbb{R}^n$ according to the lasso estimator – as treated in Example 2.5 – is in this section illuminated from the more abstract viewpoint of projections onto convex polytopes. From Lemma 3 in Tibshirani & Taylor (2012) the residual $y - \hat{\xi}^\lambda_{\text{lasso}}(y)$ is characterized as the orthogonal projection onto the closed convex polytope

$$C_\lambda = \{u \in \mathbb{R}^n \mid |\langle X_i, u\rangle| \leq \lambda \text{ for } i = 1, \ldots, p\}.$$

For a face F of C_λ determined by the affine constraints $\langle X_i, u\rangle = s_i\lambda$ for $i \in A(F)$ and $s \in \{-1, 1\}^{|A|}$ there is an associated cone

$$N_F = \left\{ \sum_{i \in A(F)} c_i s_i X_i \;\middle|\; c_i \geq 0 \right\},$$

which is itself a closed convex polytope. It is the normal cone of all points in the relative interior of F. The set $F + N_F$ is, as the Minkowski sum of two closed convex polytopes, again a closed convex polytope. The set $\{F + N_F \mid F \text{ is a face of } C_\lambda\}$

consists of the n-cells of a polytopal complex referred to as the normal manifold associated with C_λ in Robinson (1992). It holds that $\mathbb{R}^n = \bigcup_F F + N_F$ and that $\hat{S}^\lambda \subseteq A(F)$ on $\text{relint}(F) + N_F$. In fact, by Proposition 2.2 in Robinson (1992), the collection of sets $\{\text{relint}(F) + N_F \mid F \text{ is a face of } C_\lambda\}$ provides a partition of $\mathbb{R}^n$. By these arguments it is clear that the boundary of $F + N_F$ – being faces of a polytope – are $(n-1)$-dimensional. In terms of Example 2.5 it holds that

$$\overline{U}_A^\lambda = \bigcup_{F : A(F) = A} F + N_F,$$

which shows that $\overline{U}_A^\lambda$ is a union of $2^{|A|}$ closed convex polytopes.

A.2. **Semialgebraic sets.** Observe that for A and B subsets of $\mathbb{R}^n$ it holds that

$$\begin{aligned} \partial A &= \partial(A^c), \\ \partial(A \cup B) &\subseteq \partial A \cup \partial B, \\ \partial(A \cap B) &\subseteq \partial A \cup \partial B. \end{aligned} \tag{18}$$

Especially, the family of sets

$$\left\{ E \in \mathcal{B}(\mathbb{R}^n) \;\middle|\; \begin{array}{l} r \mapsto \mathcal{H}^{n-1}\left(\partial E \cap B(0,r)\right) \\ \text{is polynomially bounded} \end{array} \right\} \tag{19}$$

is stable under complement, finite union and finite intersection. This is a useful observation when we want to verify Assumption 2.2(c).

The following Lemma shows that *semialgebraic sets* belong to the family given by (19). A semialgebraic set is finite union of finite intersections of sets of the form $(P = 0)$ and $(Q > 0)$, where P and Q are polynomials. A multivariate polynomial is of the form (using multi-index notation)

$$P(x) = \sum_{\alpha \in A} a_\alpha x^\alpha, \quad a_\alpha \in \mathbb{R} \text{ for each } \alpha \in A,$$

with $A \subseteq \mathbb{N}^n$ finite.

Lemma A.1. *If E is semialgebraic then $r \mapsto \mathcal{H}^{n-1}\left(\partial E \cap B(0,r)\right)$ is polynomially bounded.*

Proof. By the stability under finite set operations of the family given by (19) it suffices to show that $r \mapsto \mathcal{H}^{n-1}\left((P=0) \cap B(0,r)\right)$ is polynomially bounded for any nonzero polynomial P. But this follows from Corollary 1 in Loi & Phien (2014), which implies that

$$\mathcal{H}^{n-1}((P=0) \cap B(0,r)) \leq \frac{\deg(P)\pi^{\frac{n+1}{2}}}{\Gamma\left(\frac{n}{2}\right)} r^{n-1}$$

for any nonzero polynomial P with $\deg(P) = \max_{a_\alpha \neq 0} |\alpha|$ denoting the degree of P. □

A.3. **Proof of Theorem 2.4.** The following Lemma characterizes the outer unit normal vectors η_i for $i = 1, \ldots, N$.

Lemma A.2. *Under Assumption 2.2 the following holds:*

(1) $\partial U_i \subseteq \bigcup_{j \neq i} \partial U_j$, for each $i = 1, ..., N$.
(2) $\eta_i = -\eta_j$ $\mathcal{H}^{n-1}$ a.e. on $\partial U_i \cap \partial U_j$ with $i \neq j$.
(3) $\eta_i = 0$ $\mathcal{H}^{n-1}$ a.e. on $\partial U_i \cap \partial U_j \cap \partial U_k$ with i, j, k distinct.

Proof. We first prove the following claim:

$$U_i^c = \bigcup_{j \neq i} \overline{U}_j, \text{ for each } i = 1, ..., N. \tag{20}$$

The inclusion $\supseteq$ is clear, since $(U_i)_i$ are open and disjoint. As for the other inclusion; if $x \notin \overline{U}_j$ for all $j \neq i$, then (since $\mathbb{R}^n = \bigcup_j \overline{U}_j$) we have $x \in \overline{U}_i$. So

$$\left(\bigcup_{j \neq i} \overline{U}_j \right)^c = \bigcap_{j \neq i} \overline{U}_j^c \subseteq \overline{U}_i$$

the left hand side is open and we conclude

$$\left(\bigcup_{j \neq i} \overline{U}_j \right)^c \subseteq \text{int}(\overline{U}_i) = U_i$$

thus proving the other inclusion of (20).

Now (1) follows from (18);

$$\partial U_i = \partial U_i^c = \partial \left(\bigcup_{j \neq i} \overline{U}_j \right) \subseteq \bigcup_{j \neq i} \partial \overline{U}_j = \bigcup_{j \neq i} \partial U_j, \tag{21}$$

where the last equality is due to

$$\partial \overline{U}_j = \overline{U}_j \setminus \text{int}(\overline{U}_j) = \overline{U}_j \setminus U_j = \partial U_j.$$

For the two other properties, note that the unit outer normal η_i on ∂U_i vanishes outside *the measure theoretic boundary* $\partial_* U_i$, see Definition 5.8 in Evans & Gariepy (1992). One can verify directly from the definition that (18) also hold for the measure theoretic boundary, hence (1) also holds for the measure theoretic boundary. Moreover, we need a third type of boundary, namely *the reduced boundary* $\partial^* U_i$, see Definition 5.7 in Evans & Gariepy (1992). Generally, the following inclusions hold:

$$\partial^* U_i \subseteq \partial_* U_i \subseteq \partial U_i.$$

Furthermore, $\mathcal{H}^{n-1}(\partial_* U_i \setminus \partial^* U_i) = 0$ (see Lemma 5.8.1 in Evans & Gariepy (1992)). All in all, we see that (2) and (3) are true, if we can show the following claim:

$$\begin{aligned} \eta_i &= -\eta_j \text{ on } \partial^* U_i \cap \partial^* U_j \\ \partial^* U_i &\cap \partial^* U_j \cap \partial^* U_k = \emptyset \end{aligned} \tag{22}$$

holds for all i, j, k distinct.

To prove the claim, define for each i and $r > 0$ the sets

$$\begin{aligned} U_i^r(x) &= \{y \mid r(y - x) + x \in U_i\}, \\ H_i(x) &= \{y \mid \langle \eta_i, y - x \rangle \leq 0\}. \end{aligned}$$

Note that $\{U_i^r(x)\}_i$ are still disjoint. By Theorem 5.7.1 in Evans & Gariepy (1992)

$$1_{U_i^r(x)} \xrightarrow{r \to 0} 1_{H_i(x)} \text{ in } L^1_{\text{loc}}(\mathbb{R}^n) \text{ for all } x \in \partial^* U_i.$$

Therefore, if there existed $x \in \partial^* U_i \cap \partial^* U_j \cap \partial^* U_k$ for i, j, k distinct, then

$$1_{U_i^r(x) \cup U_j^r(x) \cup U_k^r(x)} \xrightarrow{r \to 0} 1_{H_i(x)} + 1_{H_j(x)} + 1_{H_k(x)} \text{ in } L^1_{\text{loc}}(\mathbb{R}^n), \tag{23}$$

which is impossible as the right hand side is not Lebesgue a.e. an indicator. For the same reason one can deduce that $\eta_i = -\eta_j$ must hold for $x \in \partial^* U_i \cap \partial^* U_j$, thus we have proved the claim (22). □

Proof of Theorem 2.4. For $i = 1, ..., N$ Gauss-Green's formula (see Theorem 5.8.1 in Evans & Gariepy (1992) and Theorem 4.5.6 in Federer (1969)) gives that

$$\int_{U_i} \operatorname{div}(f)\ dm = \int_{\partial U_i} \langle f, \eta_i \rangle\ d\mathcal{H}^{n-1} \tag{24}$$

for all Lipschitz continuous vector fields f with compact support. Here η_i denotes the outer unit normal of ∂U_i, which is well defined and nonzero on a subset of ∂U_i and zero everywhere else by definition.

Let $(g_r)_r$ be a sequence of smooth functions with

$$g_r(x) = \begin{cases} 1 & \text{if } x \in B(0,r) \\ 0 & \text{if } x \notin B(0,r+1) \end{cases}$$

and $(g_r)_r$ and $(Dg_r)_r$ uniformly bounded. Since $\hat{\xi}_i$ is Lipschitz continuous on $\overline{U}_i \cap B(0, r+1)$ Kirzbraun's theorem ensures that $\hat{\xi}_i$ has a Lipschitz extension, $\hat{\xi}_i^r : \mathbb{R}^n \to \mathbb{R}^n$. Then $f_r = g_r \psi \hat{\xi}_i^r$ is Lipschitz continuous with compact support and $g_r \hat{\xi}_i^r = g_r \hat{\xi}$ on U_i. Then (24) applied to f_r yields

$$\int_{\partial U_i} g_r \psi \langle \hat{\xi}_i, \eta_i \rangle\ d\mathcal{H}^{n-1} = \int_{U_i} g_r \psi \operatorname{div}(\hat{\xi}_i)\ dm + \int_{U_i} \left\langle g_r D\psi + \psi D g_r, \hat{\xi}_i \right\rangle dm.$$

Due to Assumption 2.2 all integrands above are dominated by integrable functions, and by letting $r \to \infty$ Lebesgue's Dominated Convergence Theorem yields

$$\int_{\partial U_i} \psi \langle \hat{\xi}_i, \eta_i \rangle\ d\mathcal{H}^{n-1} = \int_{U_i} \psi \operatorname{div}(\hat{\xi}_i)\ dm + \int_{U_i} \langle D\psi, \hat{\xi}_i \rangle\ dm.$$

By summing over i we get

$$\operatorname{df}(\hat{\xi}) = \operatorname{df}_S(\hat{\xi}) - \sum_i \int_{\partial U_i} \psi \langle \hat{\xi}_i, \eta_i \rangle d\mathcal{H}^{n-1}. \tag{25}$$

By Lemma A.2 we see that

$$\begin{aligned} \operatorname{df}(\hat{\xi}) &= \operatorname{df}_S(\hat{\xi}) - \sum_{j \neq i} \int_{\partial U_i \cap \partial U_j} \psi \langle \hat{\xi}_i, \eta_i \rangle d\mathcal{H}^{n-1} \\ &= \operatorname{df}_S(\hat{\xi}) + \frac{1}{2} \sum_{j \neq i} \int_{\partial U_i \cap \partial U_j} \langle \hat{\xi}_j - \hat{\xi}_i, \eta_i \rangle \psi d\mathcal{H}^{n-1}. \end{aligned}$$

Since $\partial U_i \cap \partial U_j = \overline{U}_i \cap \overline{U}_j$ we have proven (4). □

References

Breiman, L. (1992), 'The little bootstrap and other methods for dimensionality selection in regression: X-fixed prediction error', *Journal of the American Statistical Association* **87**(419), 738–754.
URL: *http://www.tandfonline.com/doi/abs/10.1080/01621459.1992.10475276*

Bühlmann, P. & van de Geer, S. (2011), *Statistics for high-dimensional data*, Springer Series in Statistics, Springer, Heidelberg. Methods, theory and applications.

Efron, B. (2004), 'The estimation of prediction error: Covariance penalties and cross-validation', *Journal of the American Statistical Association* pp. 99–467.

Efron, B., Hastie, T., Johnstone, I. & Tibshirani, R. (2004), 'Least angle regression', *Ann. Statist.* **32**(2), 407–499. With discussion, and a rejoinder by the authors.
URL: *http://dx.doi.org/10.1214/009053604000000067*

Evans, L. & Gariepy, R. (1992), *Measure Theory and Fine Properties of Functions*, Studies in Advanced Mathematics, Taylor & Francis.

Federer, H. (1969), *Geometric measure theory*, Grundlehren der mathematischen Wissenschaften, Springer.

Friedman, J., Hastie, T. & Tibshirani, R. (2010), 'Regularization paths for generalized linear models via coordinate descent', *Journal of Statistical Software* **33**(1), 1–22.
URL: *http://www.jstatsoft.org/v33/i01/*

Givens, G. & Hoeting, J. (2012), *Computational Statistics*, Wiley Series in Computational Statistics, John Wiley & Sons, Hoboken.

Hansen, N. R. & Sokol, A. (2014), 'Degrees of freedom for nonlinear least squares estimation'.
URL: *http://arxiv.org/abs/1402.2997*

Hastie, T. J. & Tibshirani, R. J. (1990), *Generalized additive models*, Vol. 43 of *Monographs on Statistics and Applied Probability*, Chapman and Hall Ltd., London.

Kato, K. (2009), 'On the degrees of freedom in shrinkage estimation', *Journal of Multivariate Analysis* **100**(7), 1338 – 1352.
URL: *http://www.sciencedirect.com/science/article/pii/S0047259X08002753*

Loi, T. & Phien, P. (2014), 'Bounds of Hausdorff measures of tame sets', *Acta Mathematica Vietnamica* **39**(4), 637–647.
URL: *http://dx.doi.org/10.1007/s40306-014-0090-z*

Meinshausen, N. (2007), 'Relaxed lasso', *Computational Statistics & Data Analysis* **52**(1), 374 – 393.
URL: *http://www.sciencedirect.com/science/article/pii/S0167947306004956*

Meyer, M. & Woodroofe, M. (2000), 'On the degrees of freedom in shape-restricted regression', *Ann. Statist.* **28**(4), 1083–1104.
URL: *http://dx.doi.org/10.1214/aos/1015956708*

Reid, S., Tibshirani, R. & Friedman, J. (2015), 'A study of error variance estimation in lasso regression', *Statistica Sinica* .

Robinson, S. M. (1992), 'Normal maps induced by linear transformations', *Mathematics of Operations Research* **17**(3), 691–714.
URL: *http://www.jstor.org/stable/3689757*

Stein, C. M. (1981), 'Estimation of the mean of a multivariate normal distribution', *Ann. Statist.* **9**(6), 1135–1151.
URL: *http://dx.doi.org/10.1214/aos/1176345632*

Tibshirani, R. (1996), 'Regression shrinkage and selection via the lasso', *Journal of the Royal Statistical Society. Series B* **58**(1), 267–288.
URL: *http://www.jstor.org/stable/2346178*

Tibshirani, R. J. (2015), 'Degrees of freedom and model search', *Statistica Sinica* **25**, 1265–1296.

Tibshirani, R. J. & Taylor, J. (2012), 'Degrees of freedom in lasso problems', *Ann. Statist.* **40**(2), 1198–1232.

URL: *http://dx.doi.org/10.1214/12-AOS1003*

Vincent, M., Perell, K., Nielsen, F., Daugaard, G. & Hansen, N. (2014), 'Modeling tissue contamination to improve molecular identification of the primary tumor site of metastases', *Bioinformatics* **30**(10), 1417–1423.

Ye, J. (1998), 'On measuring and correcting the effects of data mining and model selection', *J. Amer. Statist. Assoc.* **93**(441), 120–131.
URL: *http://dx.doi.org/10.2307/2669609*

Zou, H., Hastie, T. & Tibshirani, R. (2007), 'On the degrees of freedom of the lasso', *Ann. Statist.* **35**(5), 2173–2192.
URL: *http://dx.doi.org/10.1214/009053607000000127*

Department of Mathematical Sciences, University of Copenhagen, Universitetsparken 5, 2100 Copenhagen Ø, Denmark

E-mail address, Corresponding author: frm@math.ku.dk

E-mail address: Niels.R.Hansen@math.ku.dk

arXiv:1601.03524v2 [math.ST] 10 Jun 2016

SUPPLEMENTARY MATERIAL: DEGREES OF FREEDOM FOR PIECEWISE LIPSCHITZ ESTIMATORS

FREDERIK RIIS MIKKELSEN AND NIELS RICHARD HANSEN

1. Tables used in heatmaps

The tables below provide the exact numbers used in Figures 4 and 5, i.e., the differences of the average risk estimates for the estimators $\hat{\xi}^{\hat{\lambda}_{\mathrm{df_S}}}_{\mathrm{d.lasso}}$ and $\hat{\xi}^{\hat{\lambda}_{\mathrm{df}}}_{\mathrm{d.lasso}}$ from the simulation study, scaled down by $n\sigma^2$.

			Orthogonal Design									
			σ known					σ estimated				
γ	n	α	0.1	0.2	0.5	1	2	0.1	0.2	0.5	1	2
1	100	0.5	0.00	0.00	0.05	0.19	0.50	-22.69	-4.93	-0.34	0.23	0.29
		0.3	0.00	0.00	-0.07	0.38	0.52	0.00	-0.03	-0.15	0.22	0.26
		0.1	0.07	0.36	0.24	0.51	0.53	0.14	0.20	0.09	0.23	0.25
		0.05	0.49	0.44	0.38	0.52	0.54	0.31	0.24	0.16	0.23	0.24
		0	0.54	0.54	0.55	0.54	0.55	0.25	0.25	0.23	0.24	0.24
	200	0.5	0.00	0.00	0.05	0.24	0.52	-22.78	-5.03	-0.36	0.27	0.37
		0.3	0.00	0.00	-0.07	0.41	0.53	0.00	-0.03	-0.16	0.29	0.35
		0.1	0.07	0.39	0.26	0.51	0.55	0.18	0.24	0.11	0.29	0.30
		0.05	0.51	0.45	0.39	0.53	0.55	0.37	0.27	0.20	0.30	0.31
		0	0.56	0.55	0.55	0.56	0.55	0.32	0.32	0.32	0.28	0.32
0.9	100	0.5	-0.02	0.05	0.41	0.52	0.55	-0.04	0.02	0.20	0.25	0.25
		0.3	-0.02	0.07	0.41	0.53	0.55	-0.04	0.03	0.20	0.24	0.24
		0.1	0.31	0.21	0.41	0.53	0.54	0.20	0.07	0.17	0.23	0.25
		0.05	0.48	0.38	0.44	0.53	0.54	0.30	0.15	0.17	0.24	0.25
		0	0.54	0.54	0.54	0.54	0.55	0.25	0.25	0.25	0.23	0.25
	200	0.5	0.28	0.33	0.49	0.55	0.56	0.15	0.17	0.26	0.31	0.29
		0.3	0.29	0.33	0.49	0.55	0.55	0.16	0.17	0.28	0.30	0.30
		0.1	0.32	0.33	0.49	0.54	0.55	0.17	0.16	0.26	0.30	0.30
		0.05	0.47	0.37	0.49	0.54	0.55	0.30	0.19	0.26	0.31	0.30
		0	0.55	0.55	0.56	0.56	0.55	0.30	0.32	0.31	0.31	0.31

			Empirical Design									
			σ known					σ estimated				
γ	n	α	0.1	0.2	0.5	1	2	0.1	0.2	0.5	1	2
1	100	0.5	-0.01	0.00	-0.03	0.03	0.13	0.08	0.01	-0.03	0.02	0.10
		0.3	-0.01	-0.02	-0.01	0.09	0.17	-0.01	-0.02	-0.02	0.05	0.15
		0.1	0.12	0.04	0.07	0.18	0.25	0.11	0.03	0.05	0.15	0.22
		0.05	0.21	0.13	0.14	0.22	0.26	0.18	0.10	0.11	0.19	0.23
		0	0.28	0.28	0.28	0.27	0.28	0.25	0.24	0.25	0.25	0.25
	200	0.5	-0.01	0.00	-0.02	0.06	0.14	0.03	0.00	-0.02	0.04	0.12
		0.3	-0.00	-0.01	0.00	0.10	0.17	-0.00	-0.01	-0.01	0.08	0.16
		0.1	0.11	0.04	0.09	0.17	0.23	0.11	0.04	0.07	0.16	0.23
		0.05	0.20	0.11	0.14	0.22	0.26	0.19	0.11	0.13	0.21	0.25
		0	0.27	0.27	0.27	0.27	0.27	0.26	0.26	0.26	0.26	0.26
0.9	100	0.5	0.00	0.02	0.12	0.21	0.25	-0.01	0.00	0.09	0.18	0.23
		0.3	0.01	0.02	0.12	0.21	0.26	-0.01	0.00	0.09	0.18	0.23
		0.1	0.06	0.05	0.14	0.23	0.26	0.05	0.03	0.11	0.19	0.23
		0.05	0.19	0.11	0.16	0.24	0.27	0.17	0.08	0.13	0.21	0.24
		0	0.27	0.28	0.28	0.27	0.27	0.24	0.25	0.25	0.24	0.25
	200	0.5	0.07	0.09	0.17	0.24	0.26	0.06	0.08	0.16	0.23	0.25
		0.3	0.07	0.09	0.18	0.23	0.26	0.06	0.08	0.16	0.23	0.25
		0.1	0.07	0.09	0.17	0.24	0.26	0.06	0.08	0.16	0.23	0.25
		0.05	0.13	0.12	0.19	0.25	0.26	0.13	0.10	0.17	0.23	0.26
		0	0.27	0.27	0.27	0.27	0.27	0.27	0.27	0.26	0.26	0.27

			Constant Correlation Design: $\rho = 0.1$									
			σ known					σ estimated				
γ	n	α	0.1	0.2	0.5	1	2	0.1	0.2	0.5	1	2
1	100	0.5	-0.00	0.01	0.04	0.00	-0.01	0.71	0.17	0.03	0.01	0.01
		0.3	0.02	0.02	0.02	-0.01	0.06	0.05	0.03	0.02	0.00	0.07
		0.1	0.18	0.15	0.08	0.15	0.28	0.15	0.12	0.06	0.12	0.24
		0.05	0.28	0.24	0.16	0.28	0.33	0.23	0.20	0.13	0.23	0.28
		0	0.34	0.35	0.35	0.35	0.34	0.31	0.30	0.32	0.31	0.31
	200	0.5	-0.00	0.01	0.06	0.05	-0.01	0.68	0.17	0.04	0.05	0.02
		0.3	0.04	0.07	0.06	0.01	0.03	0.06	0.07	0.06	0.03	0.07
		0.1	0.20	0.18	0.09	0.14	0.25	0.20	0.17	0.10	0.15	0.25
		0.05	0.31	0.26	0.19	0.24	0.32	0.29	0.25	0.19	0.24	0.31
		0	0.35	0.35	0.36	0.35	0.36	0.34	0.34	0.34	0.35	0.34
0.9	100	0.5	0.02	0.02	0.04	0.15	0.27	0.02	0.03	0.04	0.13	0.25
		0.3	0.03	0.02	0.05	0.17	0.29	0.03	0.03	0.05	0.15	0.25
		0.1	0.15	0.10	0.11	0.24	0.32	0.13	0.08	0.09	0.20	0.28
		0.05	0.28	0.21	0.18	0.30	0.34	0.23	0.16	0.14	0.25	0.29
		0	0.34	0.35	0.34	0.35	0.34	0.30	0.31	0.31	0.32	0.30
	200	0.5	0.08	0.09	0.15	0.25	0.33	0.10	0.11	0.17	0.25	0.32
		0.3	0.08	0.08	0.14	0.25	0.32	0.10	0.10	0.15	0.25	0.31
		0.1	0.15	0.13	0.17	0.27	0.33	0.14	0.13	0.18	0.26	0.32
		0.05	0.28	0.22	0.22	0.30	0.34	0.26	0.21	0.21	0.29	0.33
		0	0.35	0.36	0.36	0.35	0.36	0.34	0.34	0.35	0.34	0.34

			Correlated Design: $\rho = 0.4$									
			σ known					σ estimated				
γ	n	α	0.1	0.2	0.5	1	2	0.1	0.2	0.5	1	2
1	100	0.5	0.00	0.03	0.07	0.06	0.02	2.93	0.68	0.10	0.07	0.06
		0.3	-0.00	0.05	0.09	0.04	0.07	0.56	0.20	0.07	0.07	0.11
		0.1	0.11	0.17	0.12	0.18	0.25	0.09	0.13	0.11	0.18	0.24
		0.05	0.25	0.23	0.21	0.25	0.31	0.23	0.19	0.18	0.24	0.27
		0	0.35	0.34	0.35	0.35	0.35	0.31	0.30	0.32	0.31	0.31
	200	0.5	0.00	0.02	0.10	0.09	0.04	3.03	0.71	0.12	0.09	0.10
		0.3	-0.00	0.03	0.10	0.07	0.08	0.73	0.17	0.11	0.12	0.14
		0.1	0.12	0.23	0.16	0.19	0.29	0.11	0.22	0.17	0.23	0.31
		0.05	0.29	0.27	0.24	0.29	0.32	0.28	0.27	0.25	0.30	0.32
		0	0.35	0.35	0.35	0.35	0.35	0.34	0.35	0.34	0.34	0.34
0.9	100	0.5	0.05	0.08	0.07	0.15	0.26	0.07	0.08	0.09	0.16	0.24
		0.3	0.06	0.08	0.08	0.16	0.26	0.07	0.08	0.10	0.17	0.24
		0.1	0.16	0.14	0.15	0.23	0.30	0.13	0.11	0.14	0.21	0.27
		0.05	0.25	0.20	0.22	0.27	0.32	0.22	0.18	0.20	0.24	0.29
		0	0.35	0.35	0.34	0.35	0.35	0.31	0.32	0.30	0.30	0.30
	200	0.5	0.14	0.15	0.20	0.29	0.32	0.16	0.18	0.24	0.31	0.32
		0.3	0.14	0.15	0.20	0.29	0.32	0.16	0.19	0.24	0.31	0.32
		0.1	0.20	0.18	0.23	0.30	0.32	0.20	0.20	0.26	0.32	0.32
		0.05	0.27	0.24	0.28	0.31	0.32	0.27	0.24	0.29	0.32	0.32
		0	0.36	0.35	0.35	0.35	0.36	0.35	0.34	0.34	0.35	0.34

2. Errorbar plots of risk estimates

Errorbar plots of the risk estimates relative to $n\sigma^2$ for the estimators $\hat{\xi}^{\hat{\lambda}_{\mathrm{df_S}}}_{\mathrm{d.lasso}}$ and $\hat{\xi}^{\hat{\lambda}_{\mathrm{df}}}_{\mathrm{d.lasso}}$. The dashed lines are the relative risks for the oracle estimator.

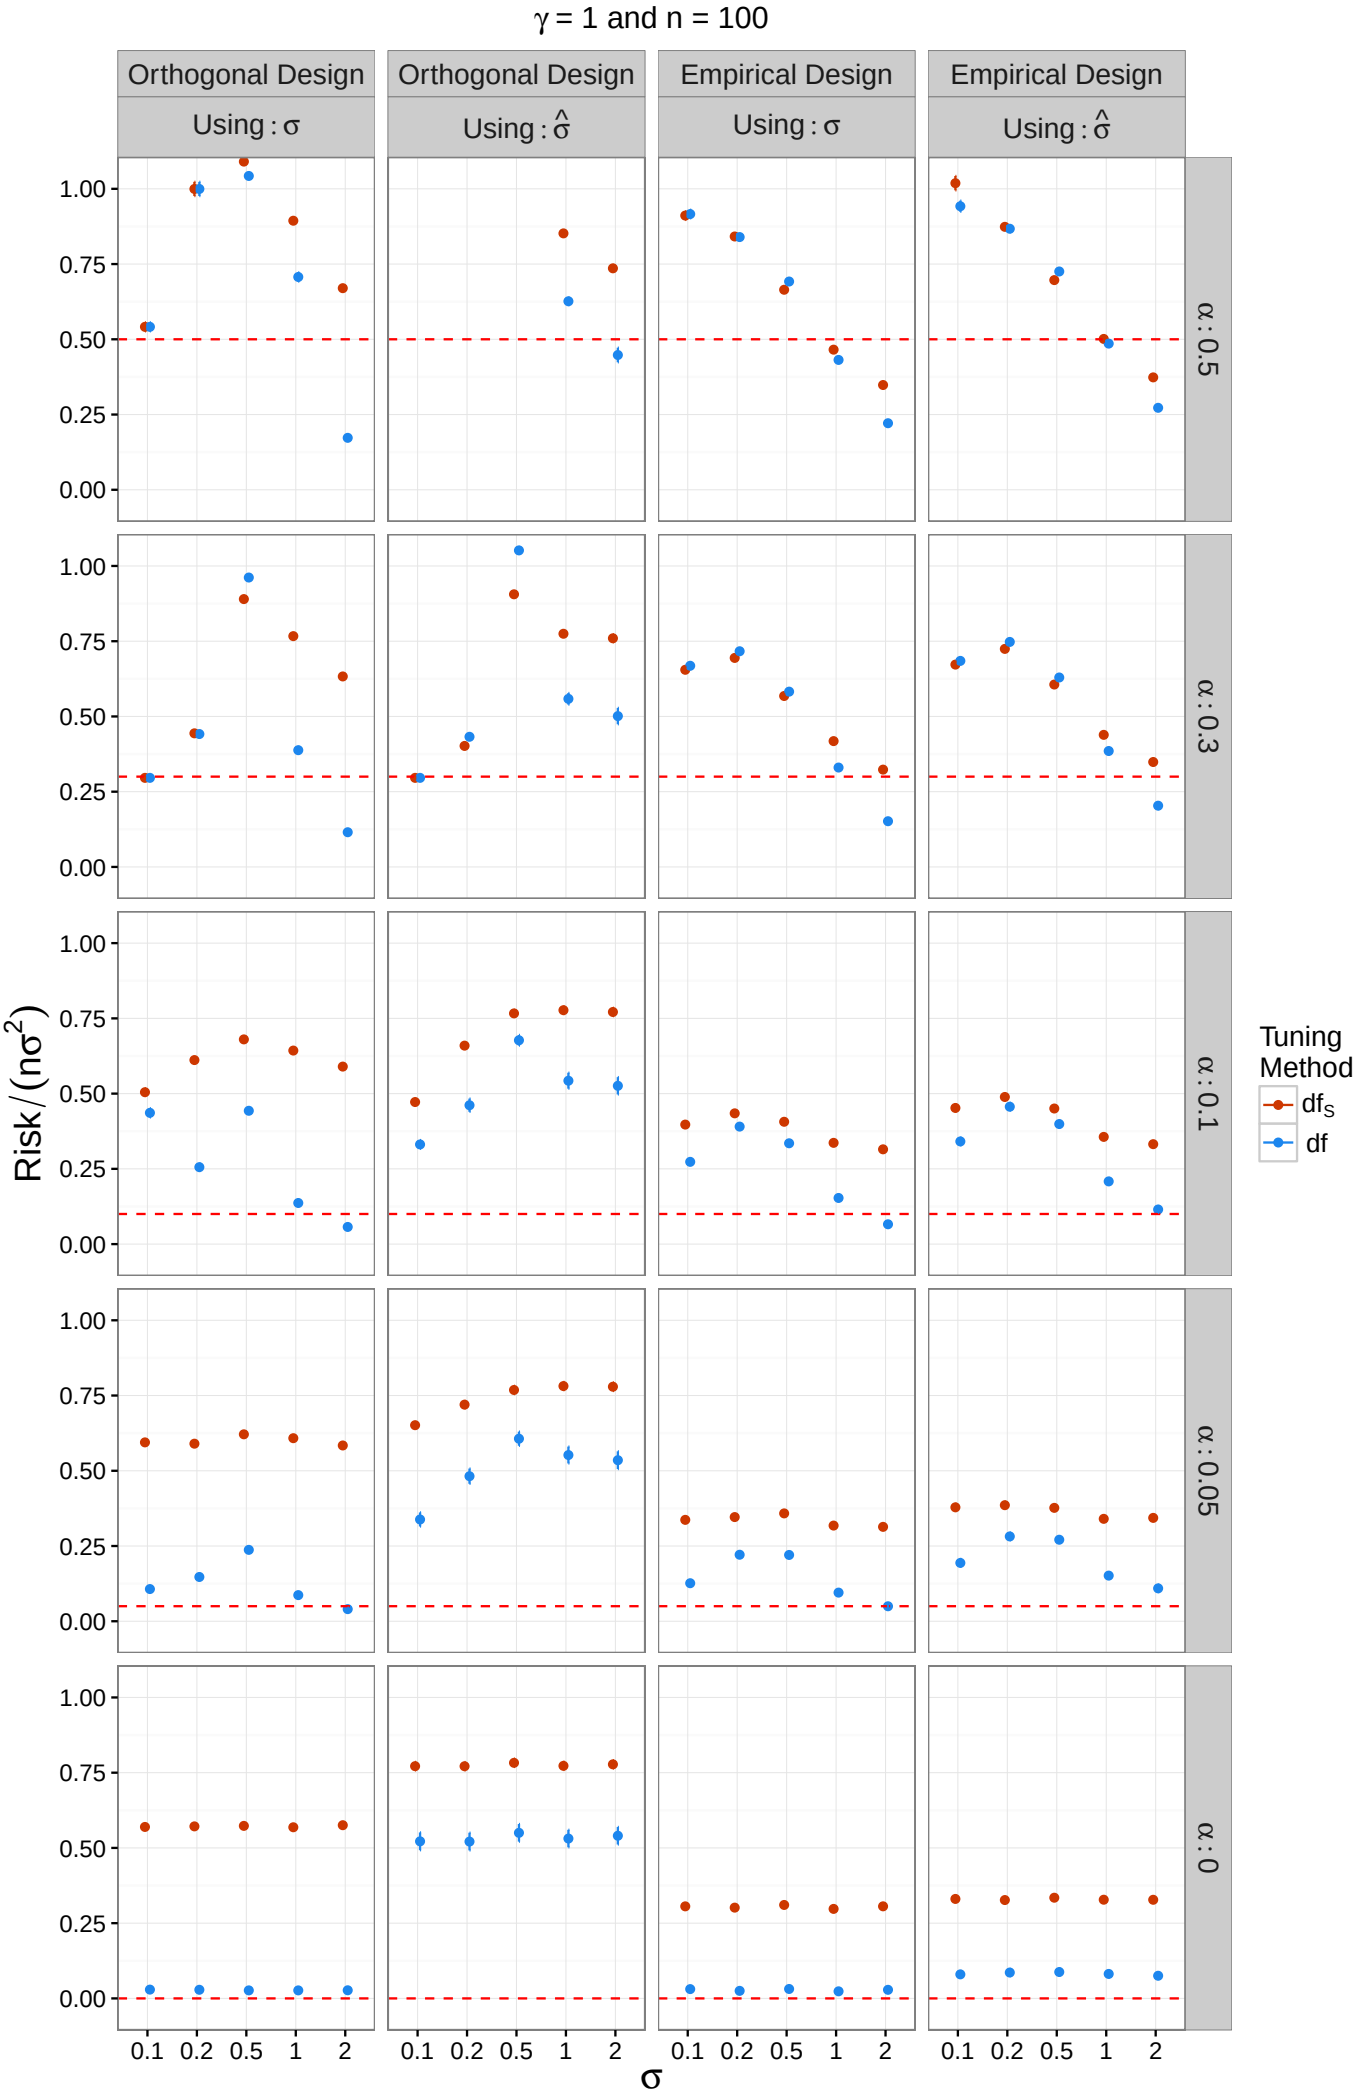
γ = 1 and n = 100
Orthogonal Design
Using : σ
Orthogonal Design
Using : σ̂
Empirical Design
Using : σ
Empirical Design
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
σ
0.1 0.2 0.5 1 2
0.00 0.25 0.50 0.75 1.00
Tuning Method
df_S
df

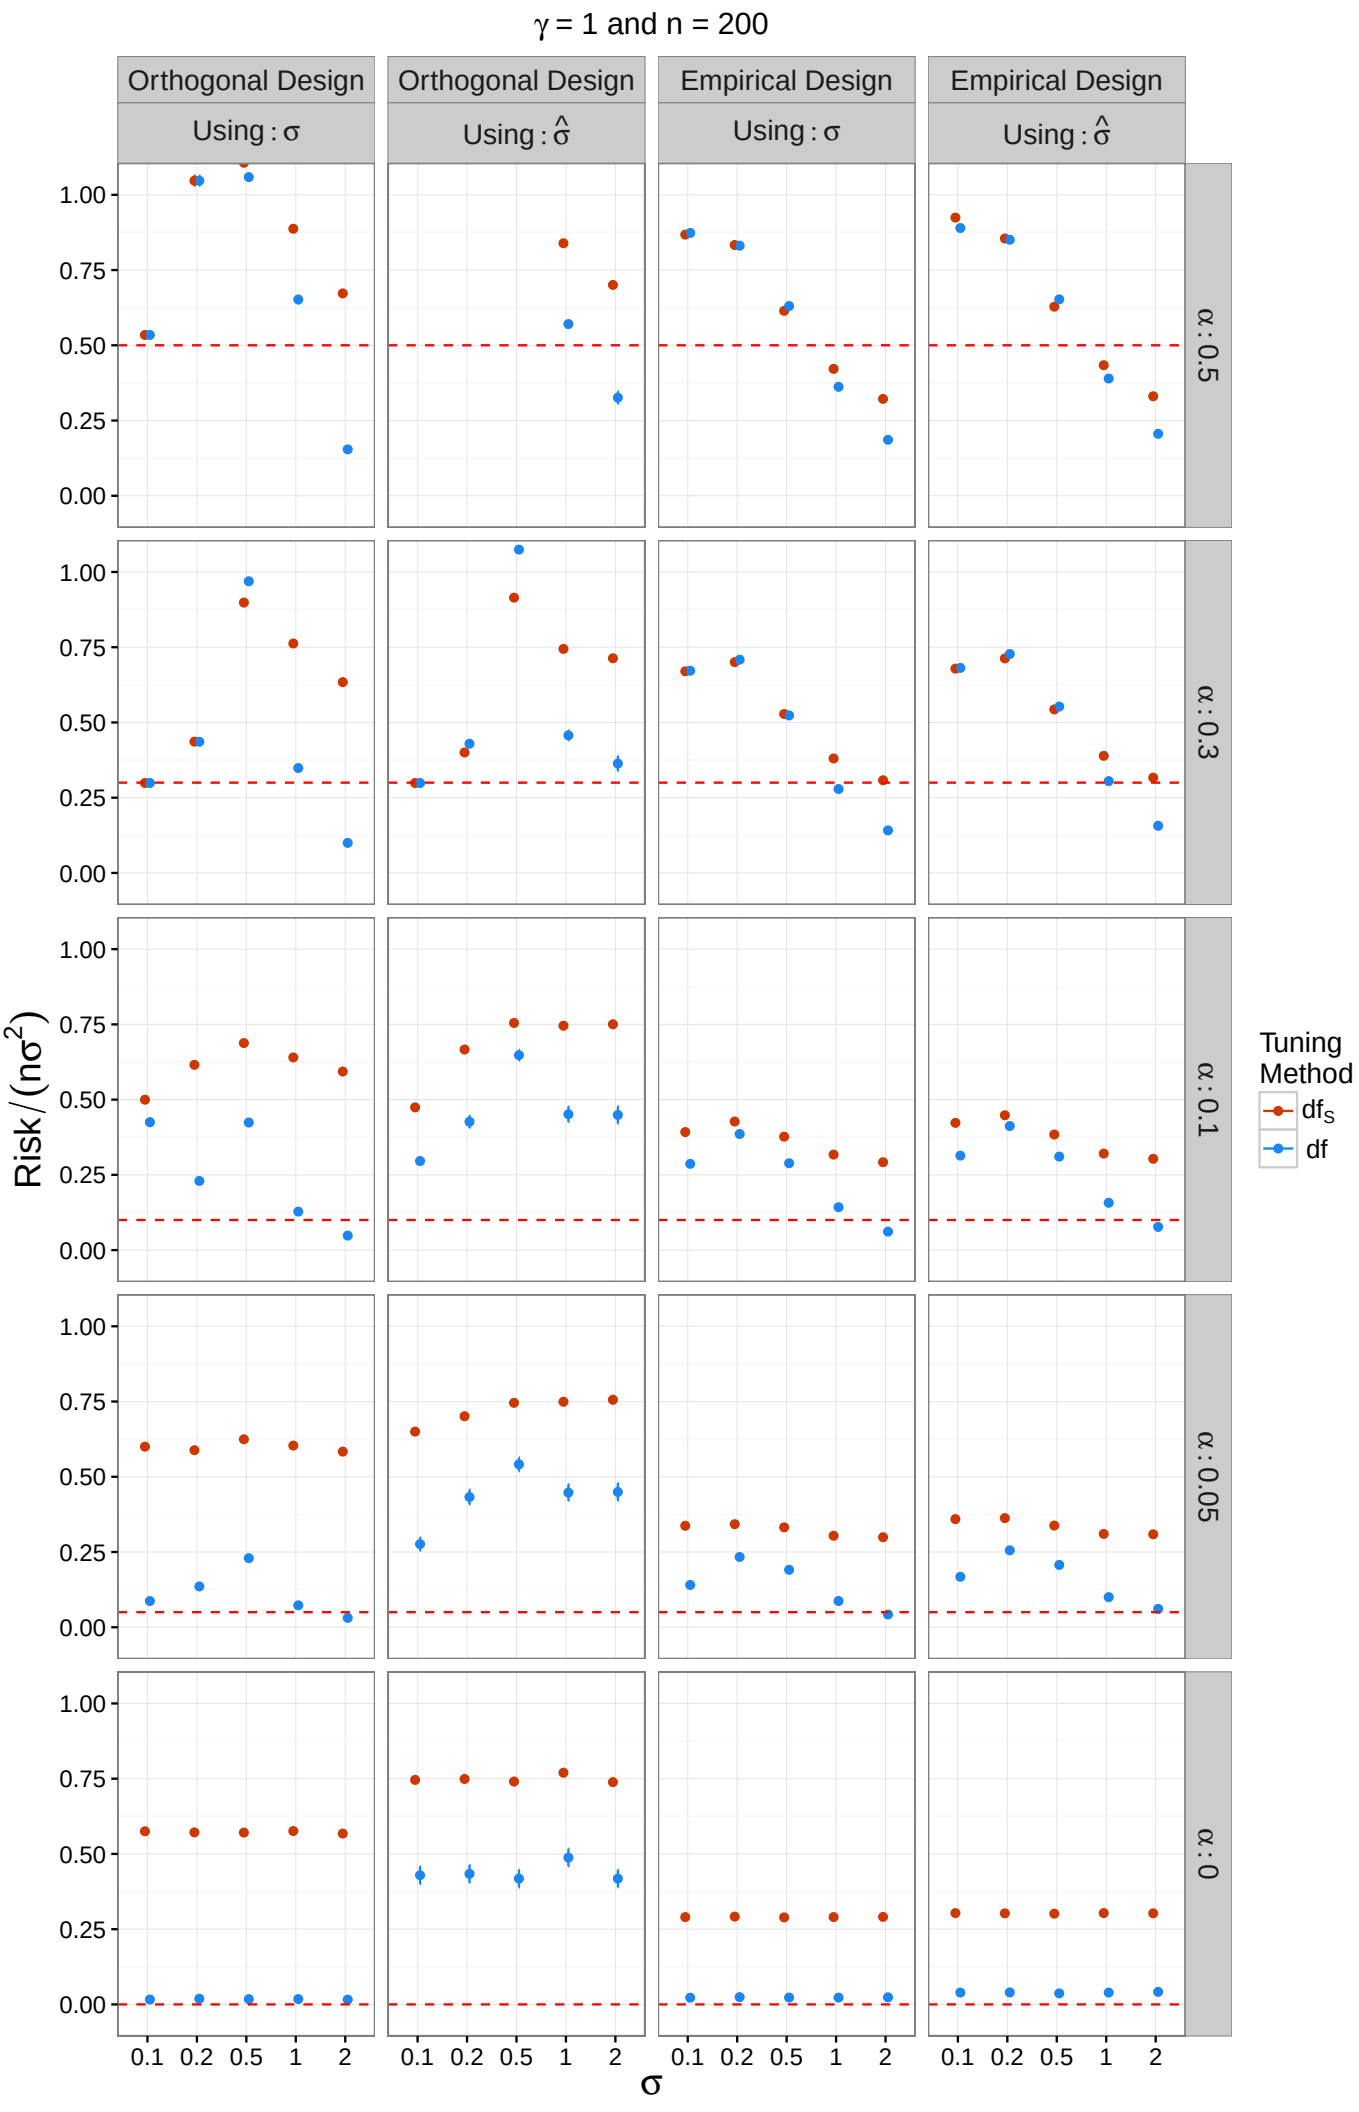
γ = 1 and n = 200
Orthogonal Design
Using : σ
Orthogonal Design
Using : σ̂
Empirical Design
Using : σ
Empirical Design
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1
0.2
0.5
1
2
σ
Tuning Method
df_S
df

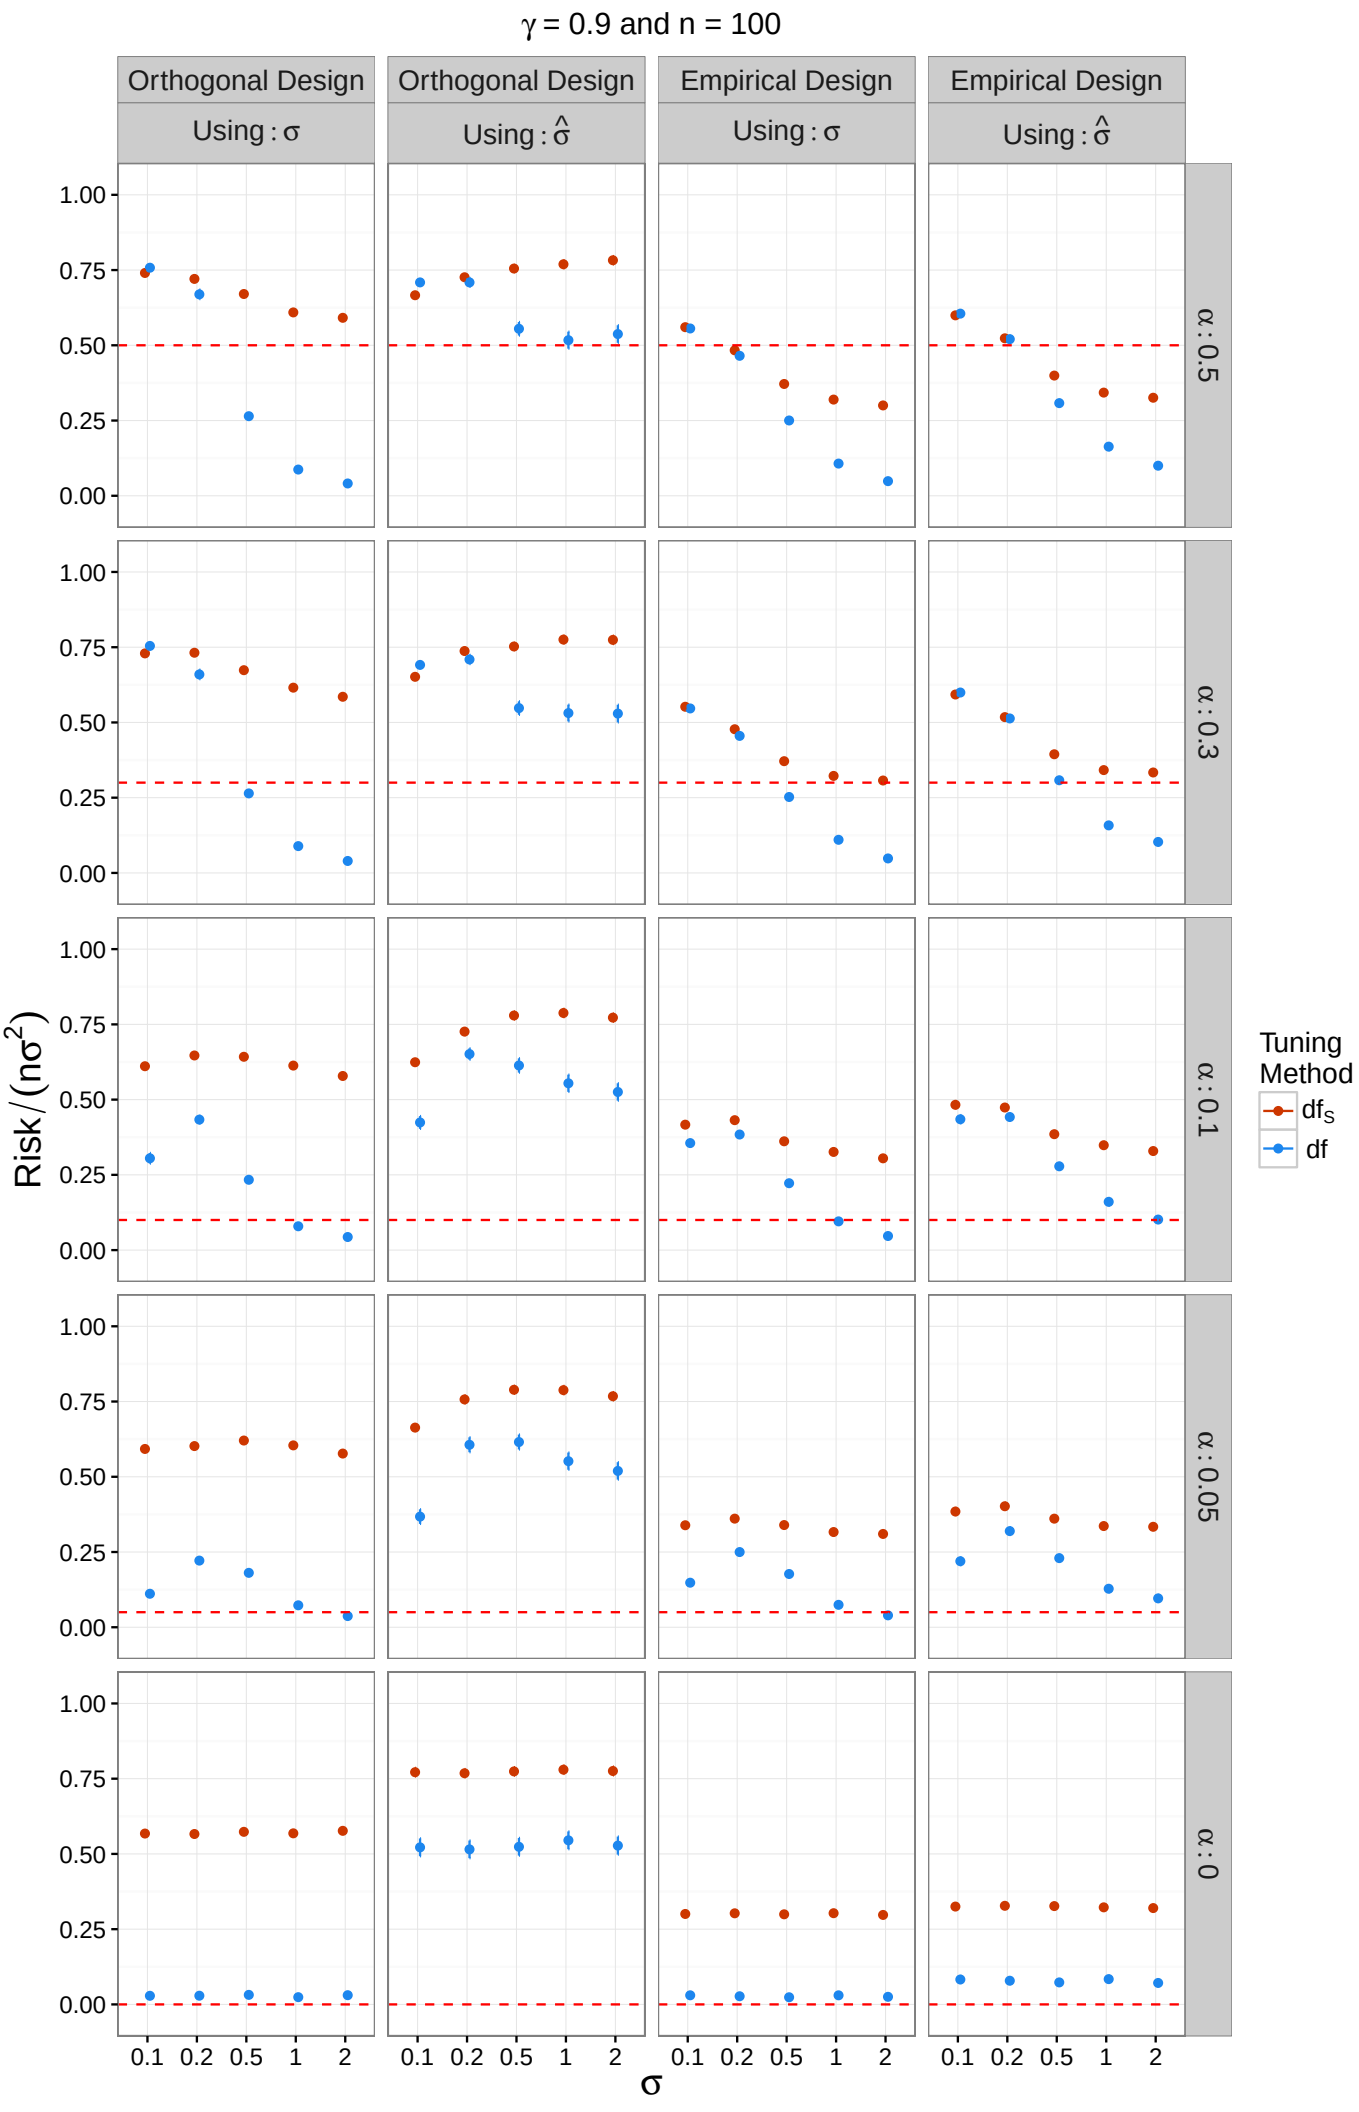
γ = 0.9 and n = 100
Orthogonal Design
Using : σ
Orthogonal Design
Using : σ̂
Empirical Design
Using : σ
Empirical Design
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1
0.2
0.5
1
2
σ
Tuning Method
df_S
df

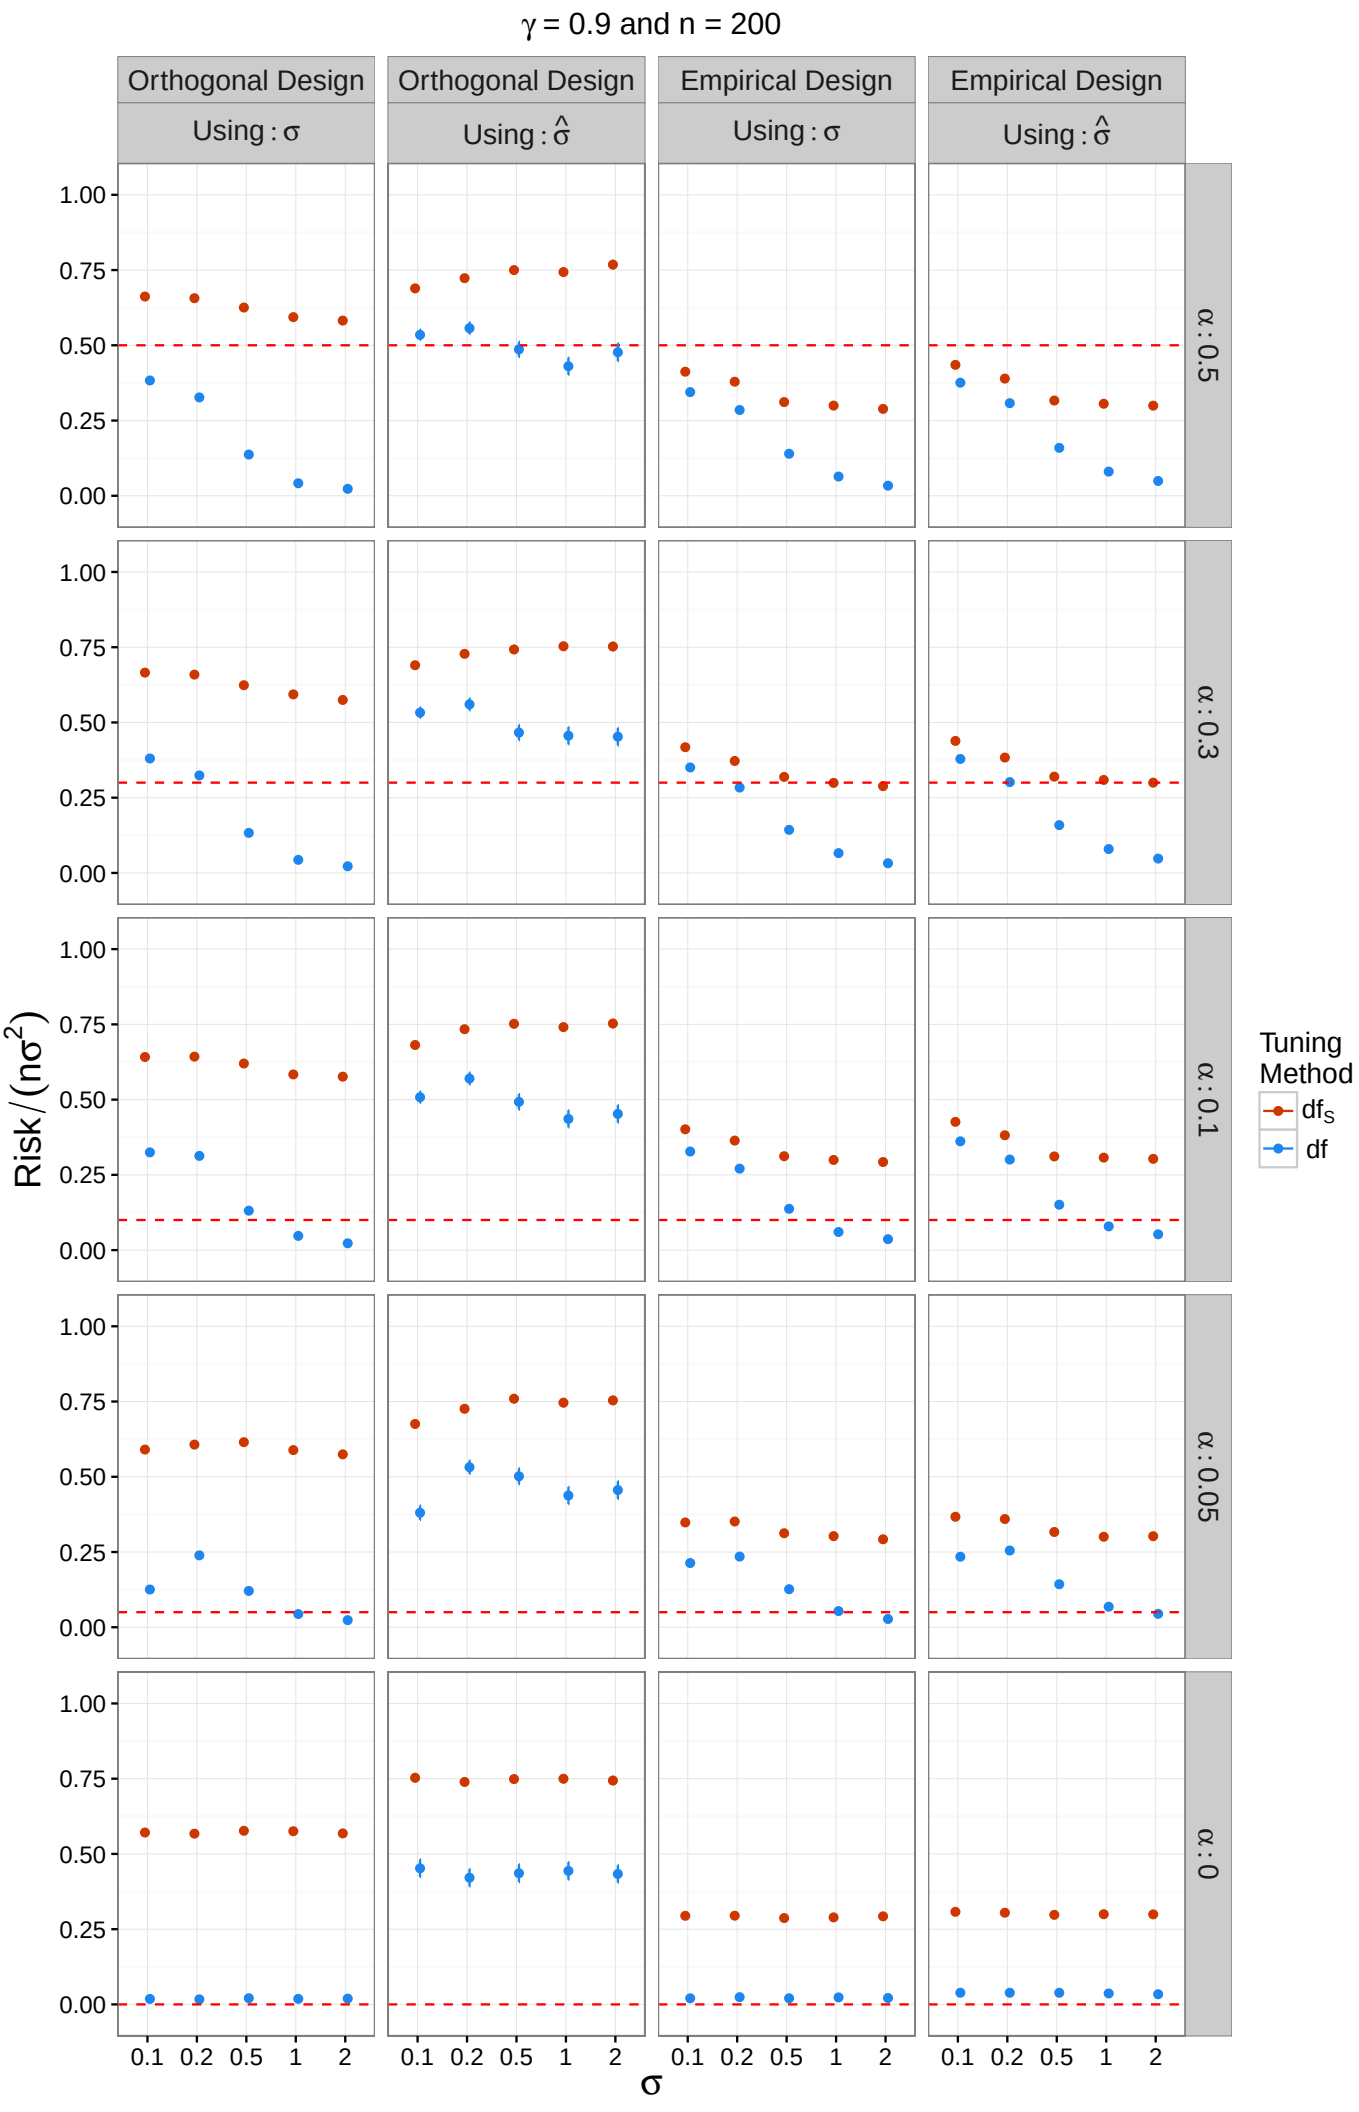
γ = 0.9 and n = 200
Orthogonal Design
Using : σ
Orthogonal Design
Using : σ̂
Empirical Design
Using : σ
Empirical Design
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1
0.2
0.5
1
2
σ
Tuning Method
df_S
df

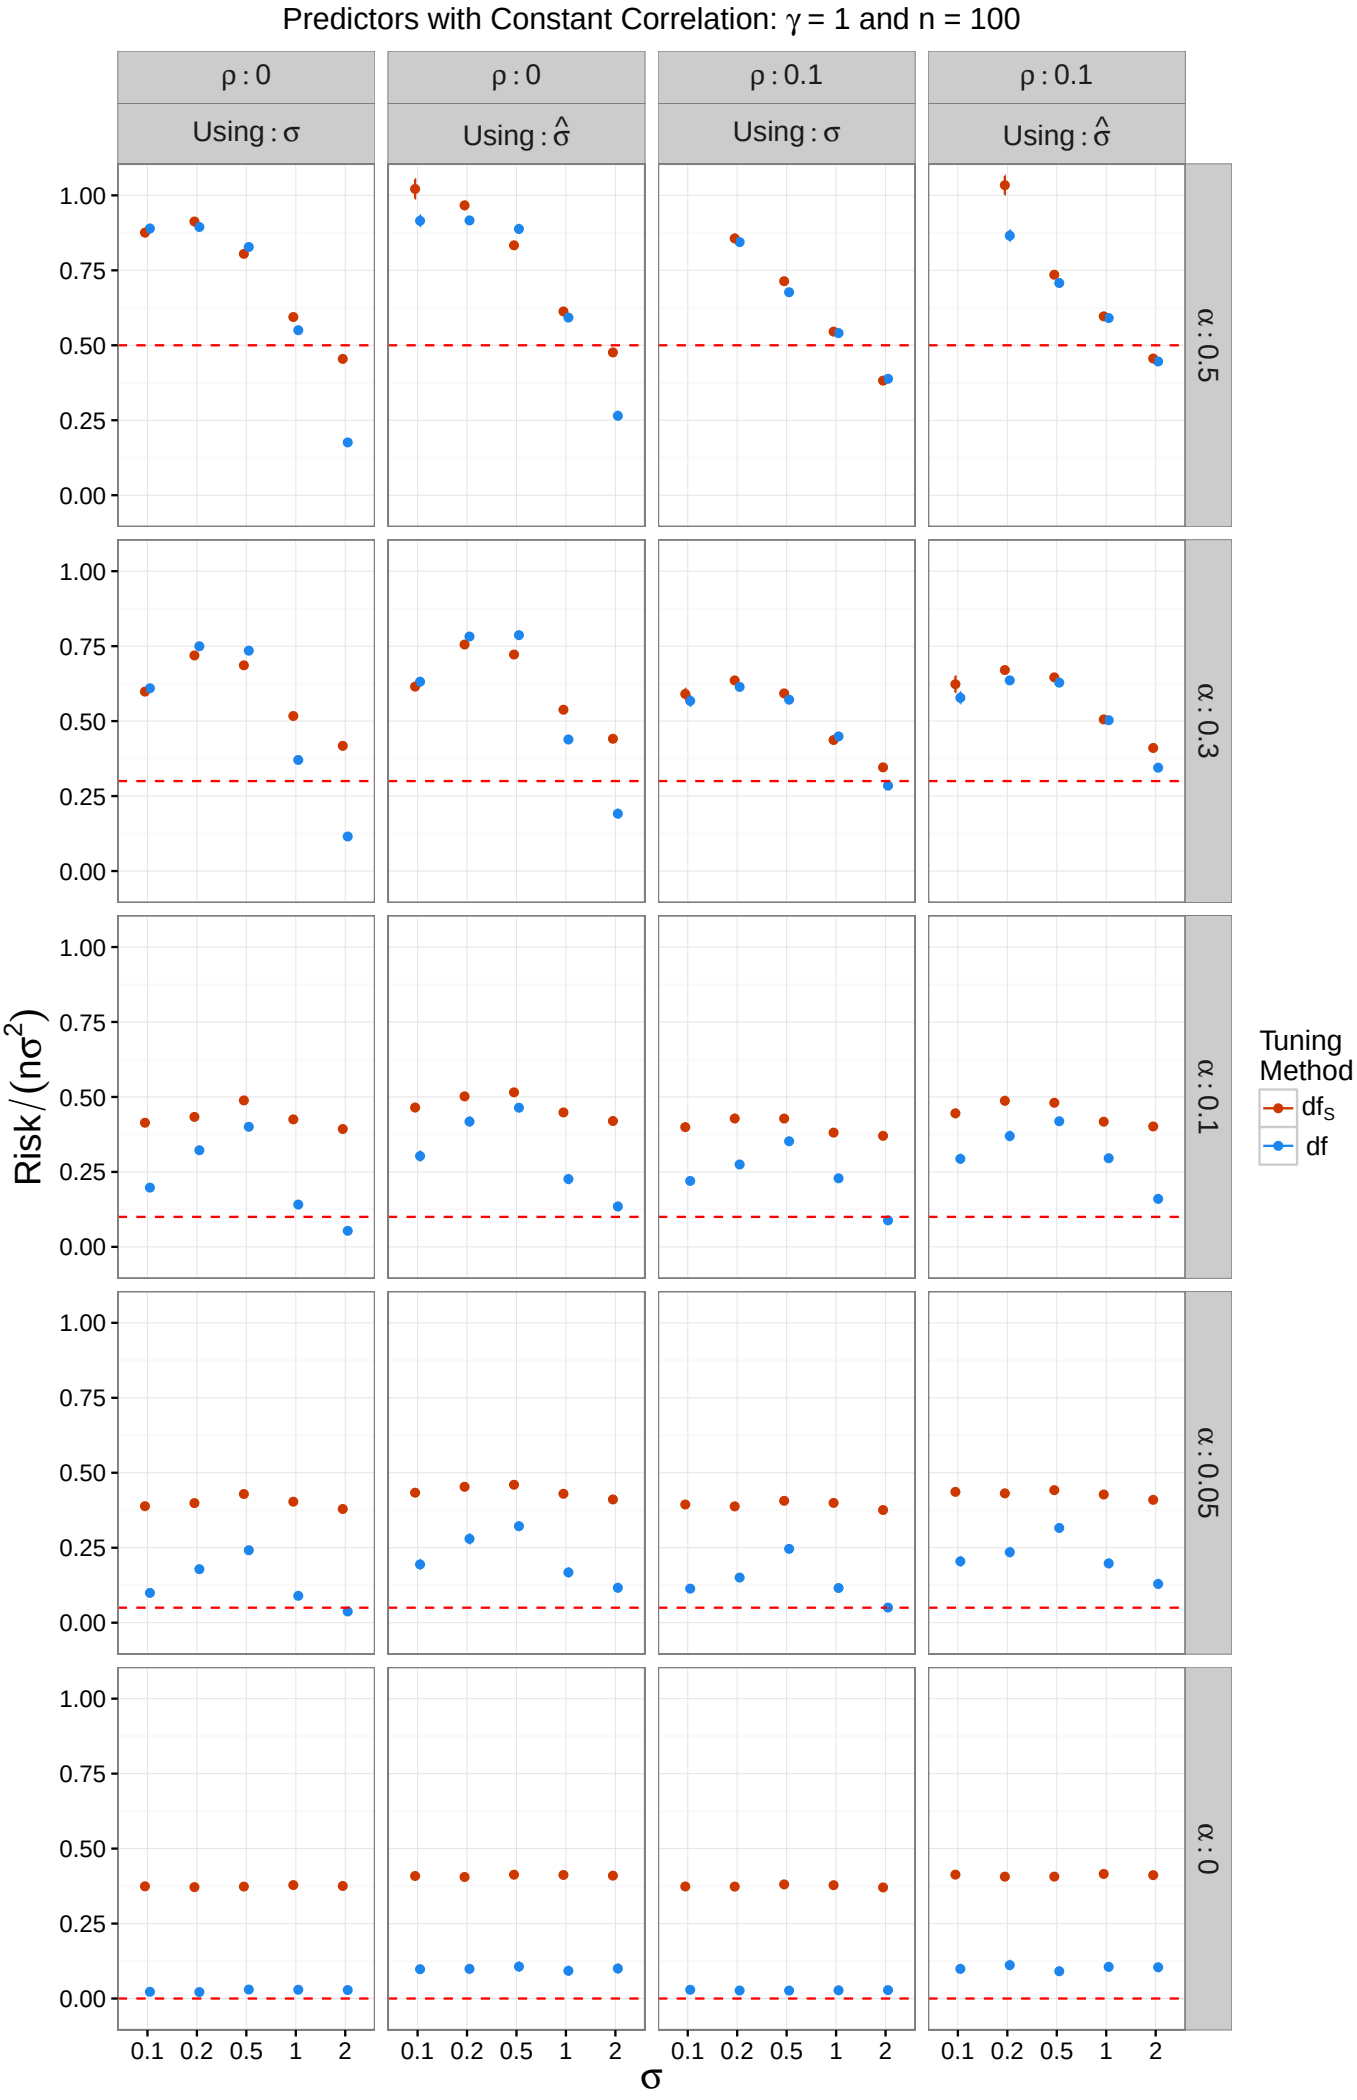
Predictors with Constant Correlation: γ = 1 and n = 100
ρ : 0
Using : σ
ρ : 0
Using : σ̂
ρ : 0.1
Using : σ
ρ : 0.1
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk / (nσ²)
1.00
0.75
0.50
0.25
0.00
0.1 0.2 0.5 1 2
σ
Tuning Method
df_S
df

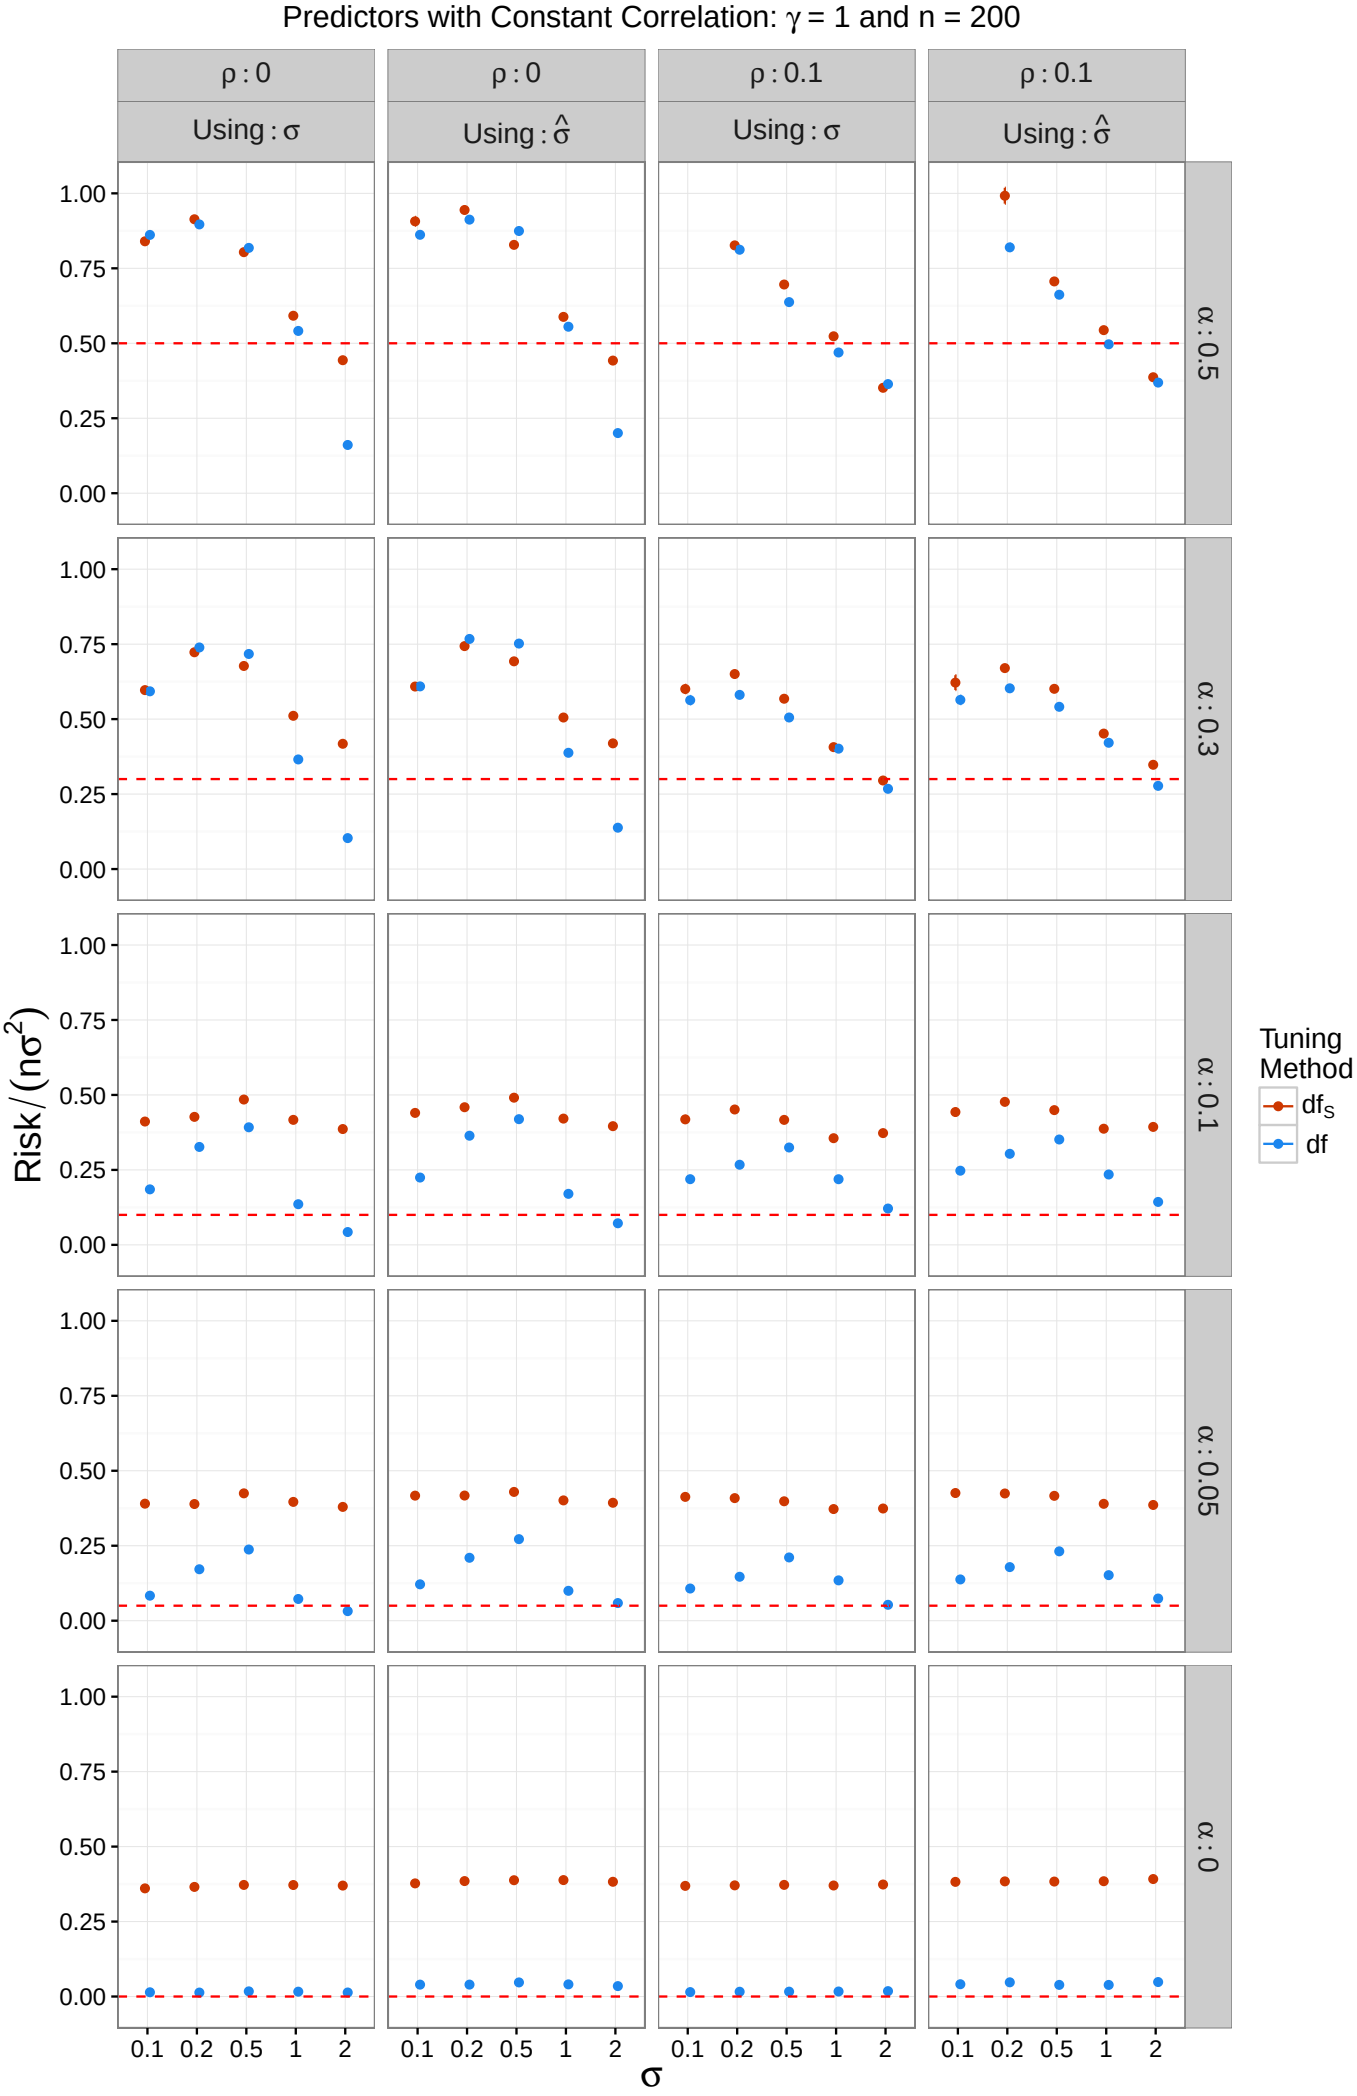
Predictors with Constant Correlation: γ = 1 and n = 200
ρ : 0
Using : σ
ρ : 0
Using : σ̂
ρ : 0.1
Using : σ
ρ : 0.1
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1 0.2 0.5 1 2
σ
Tuning Method
df_S
df

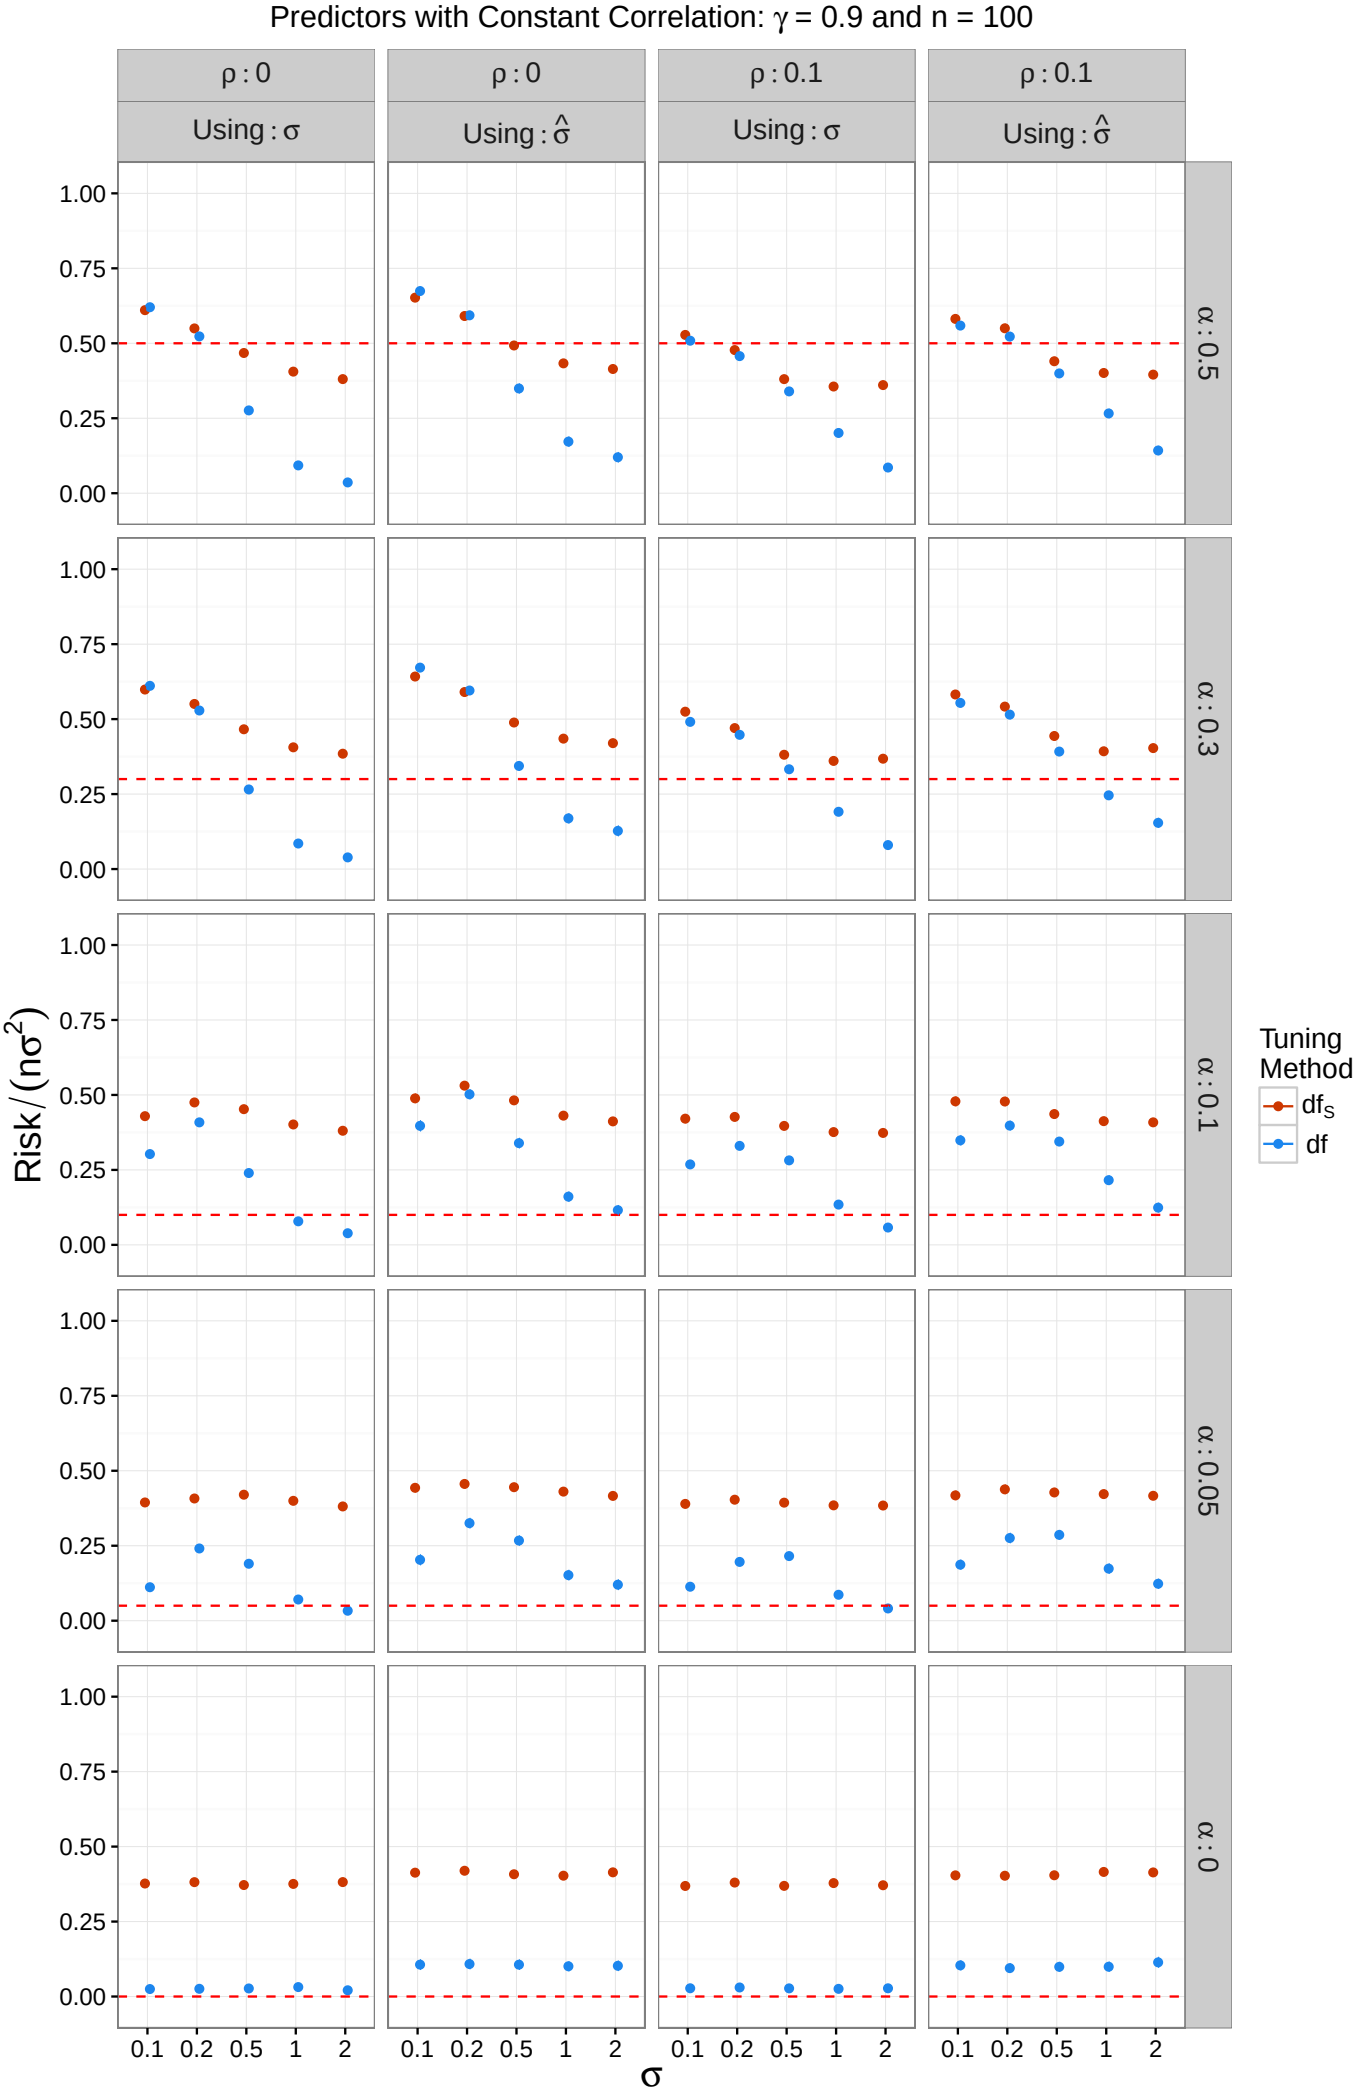
Predictors with Constant Correlation: γ = 0.9 and n = 100
ρ : 0
Using : σ
ρ : 0
Using : σ̂
ρ : 0.1
Using : σ
ρ : 0.1
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1 0.2 0.5 1 2
σ
Tuning Method
df_S
df

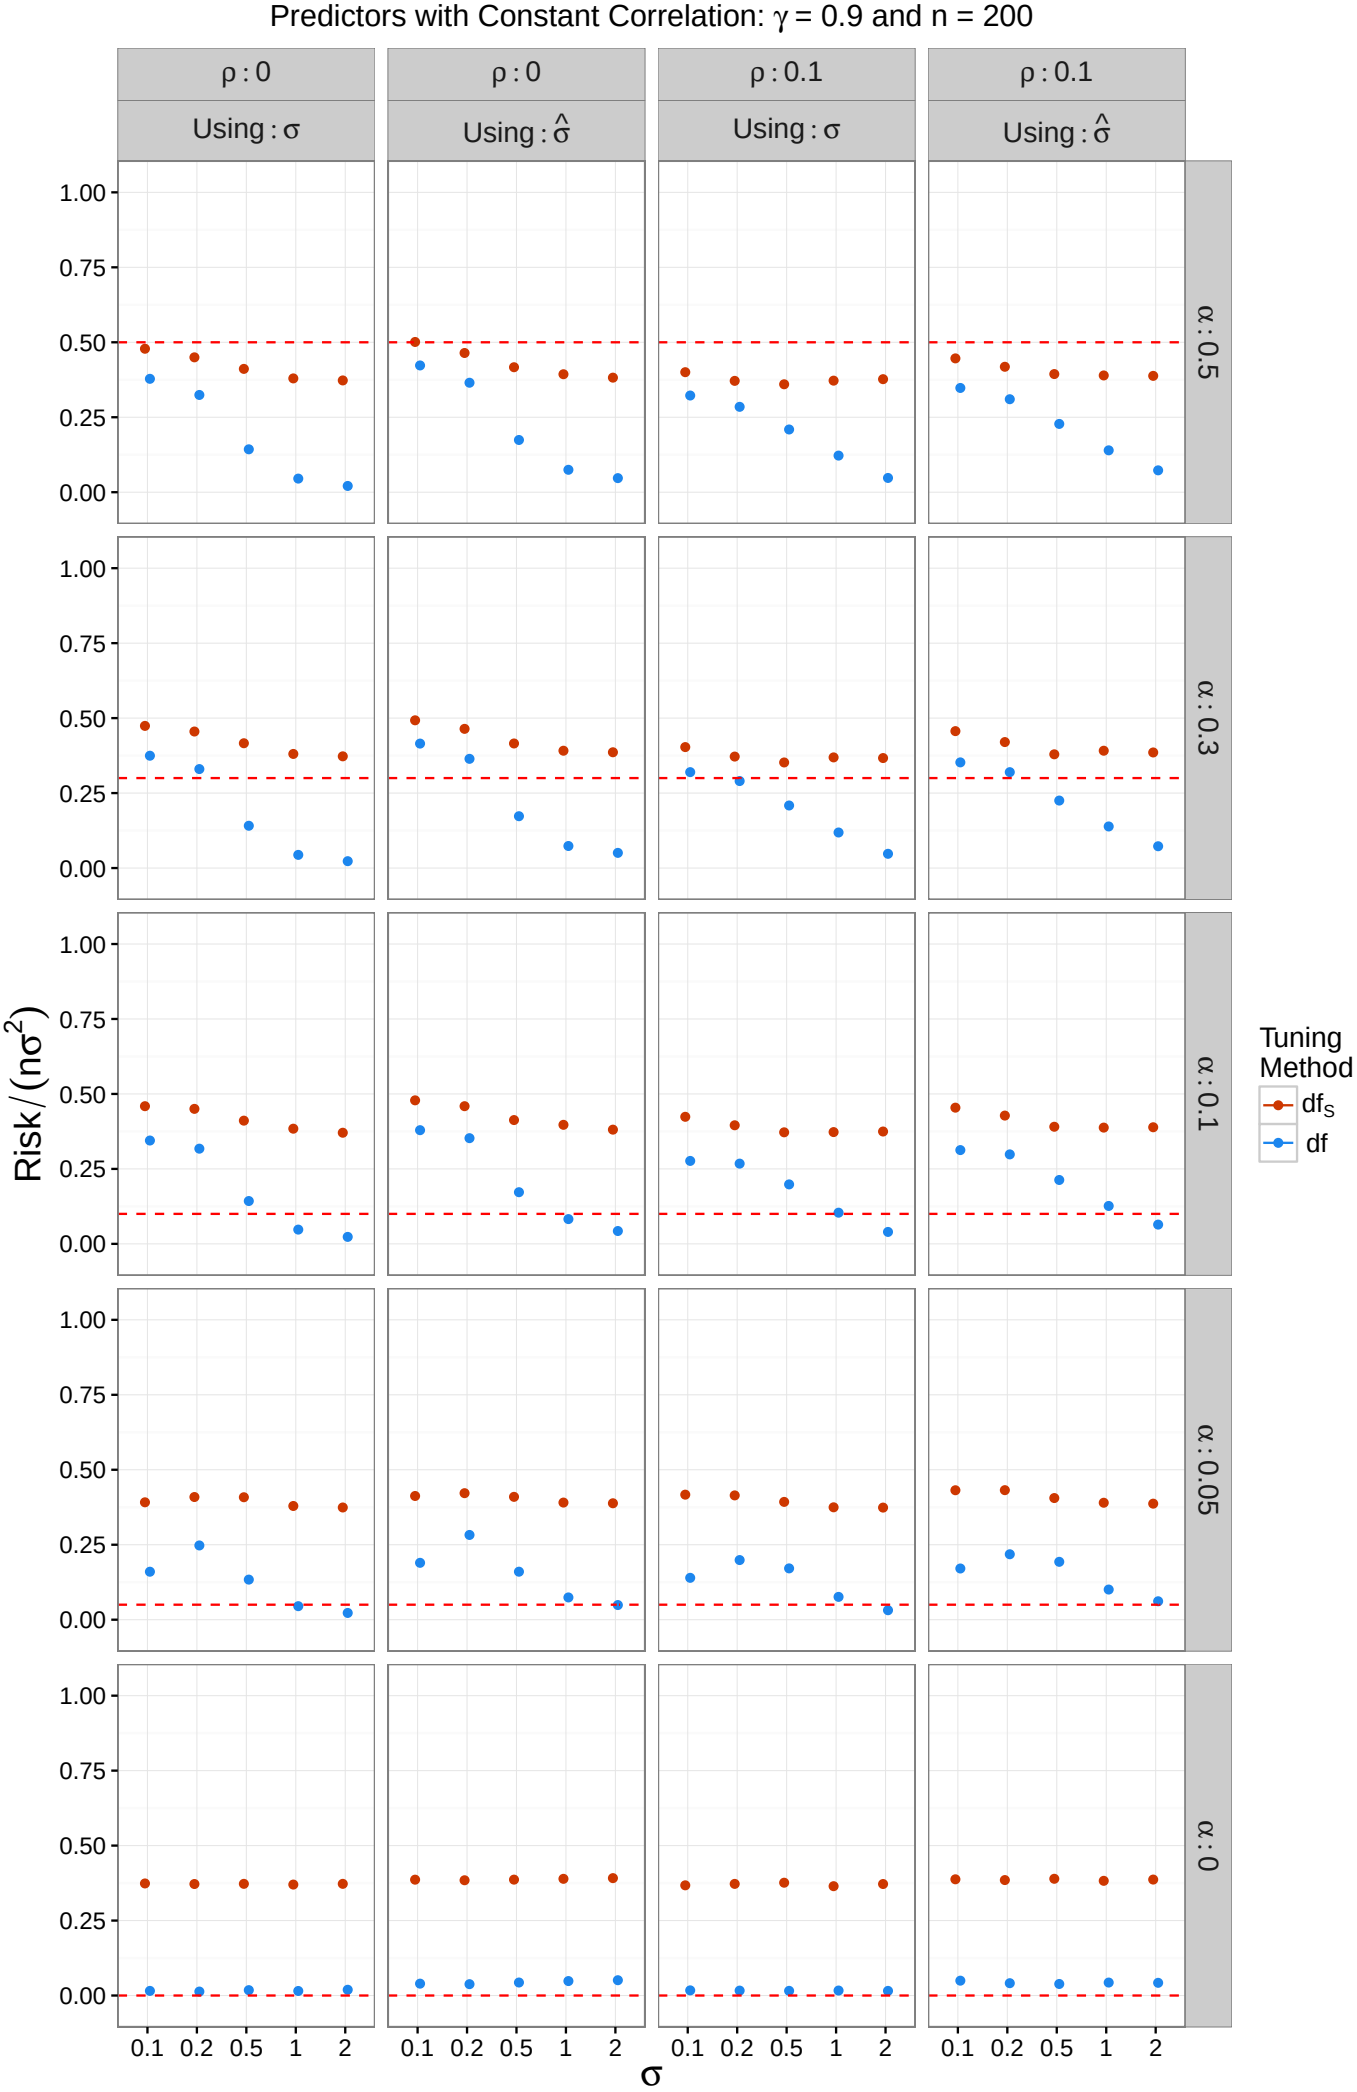
Predictors with Constant Correlation: γ = 0.9 and n = 200
ρ : 0
Using : σ
ρ : 0
Using : σ̂
ρ : 0.1
Using : σ
ρ : 0.1
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1 0.2 0.5 1 2
σ
Tuning Method
df_S
df

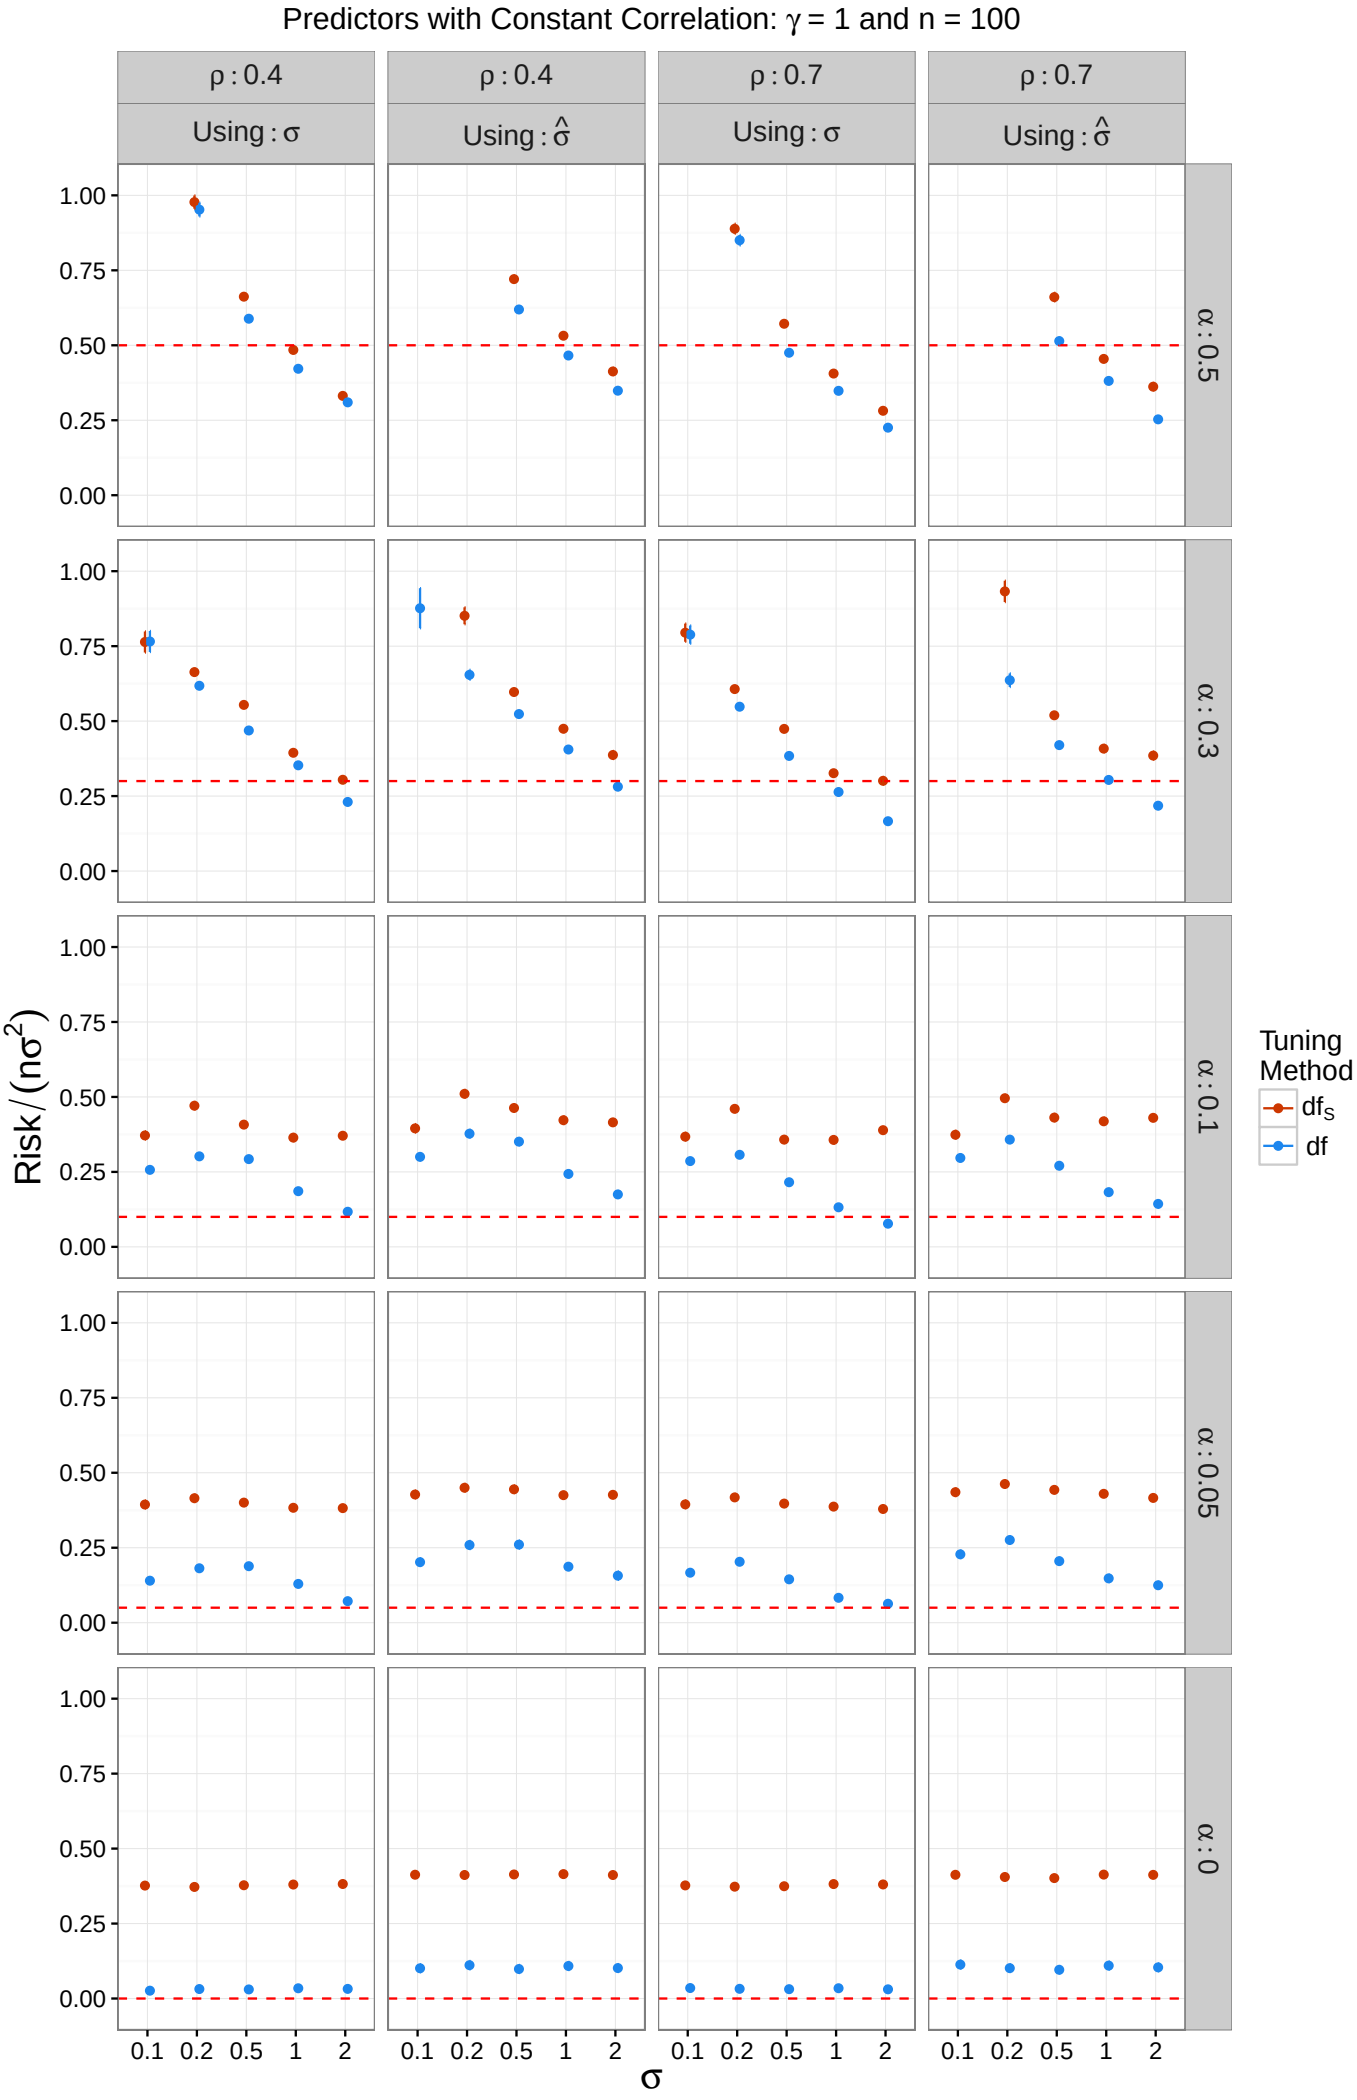
Predictors with Constant Correlation: γ = 1 and n = 100
ρ : 0.4
Using : σ
ρ : 0.4
Using : σ̂
ρ : 0.7
Using : σ
ρ : 0.7
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1
0.2
0.5
1
2
σ
Tuning Method
df_S
df

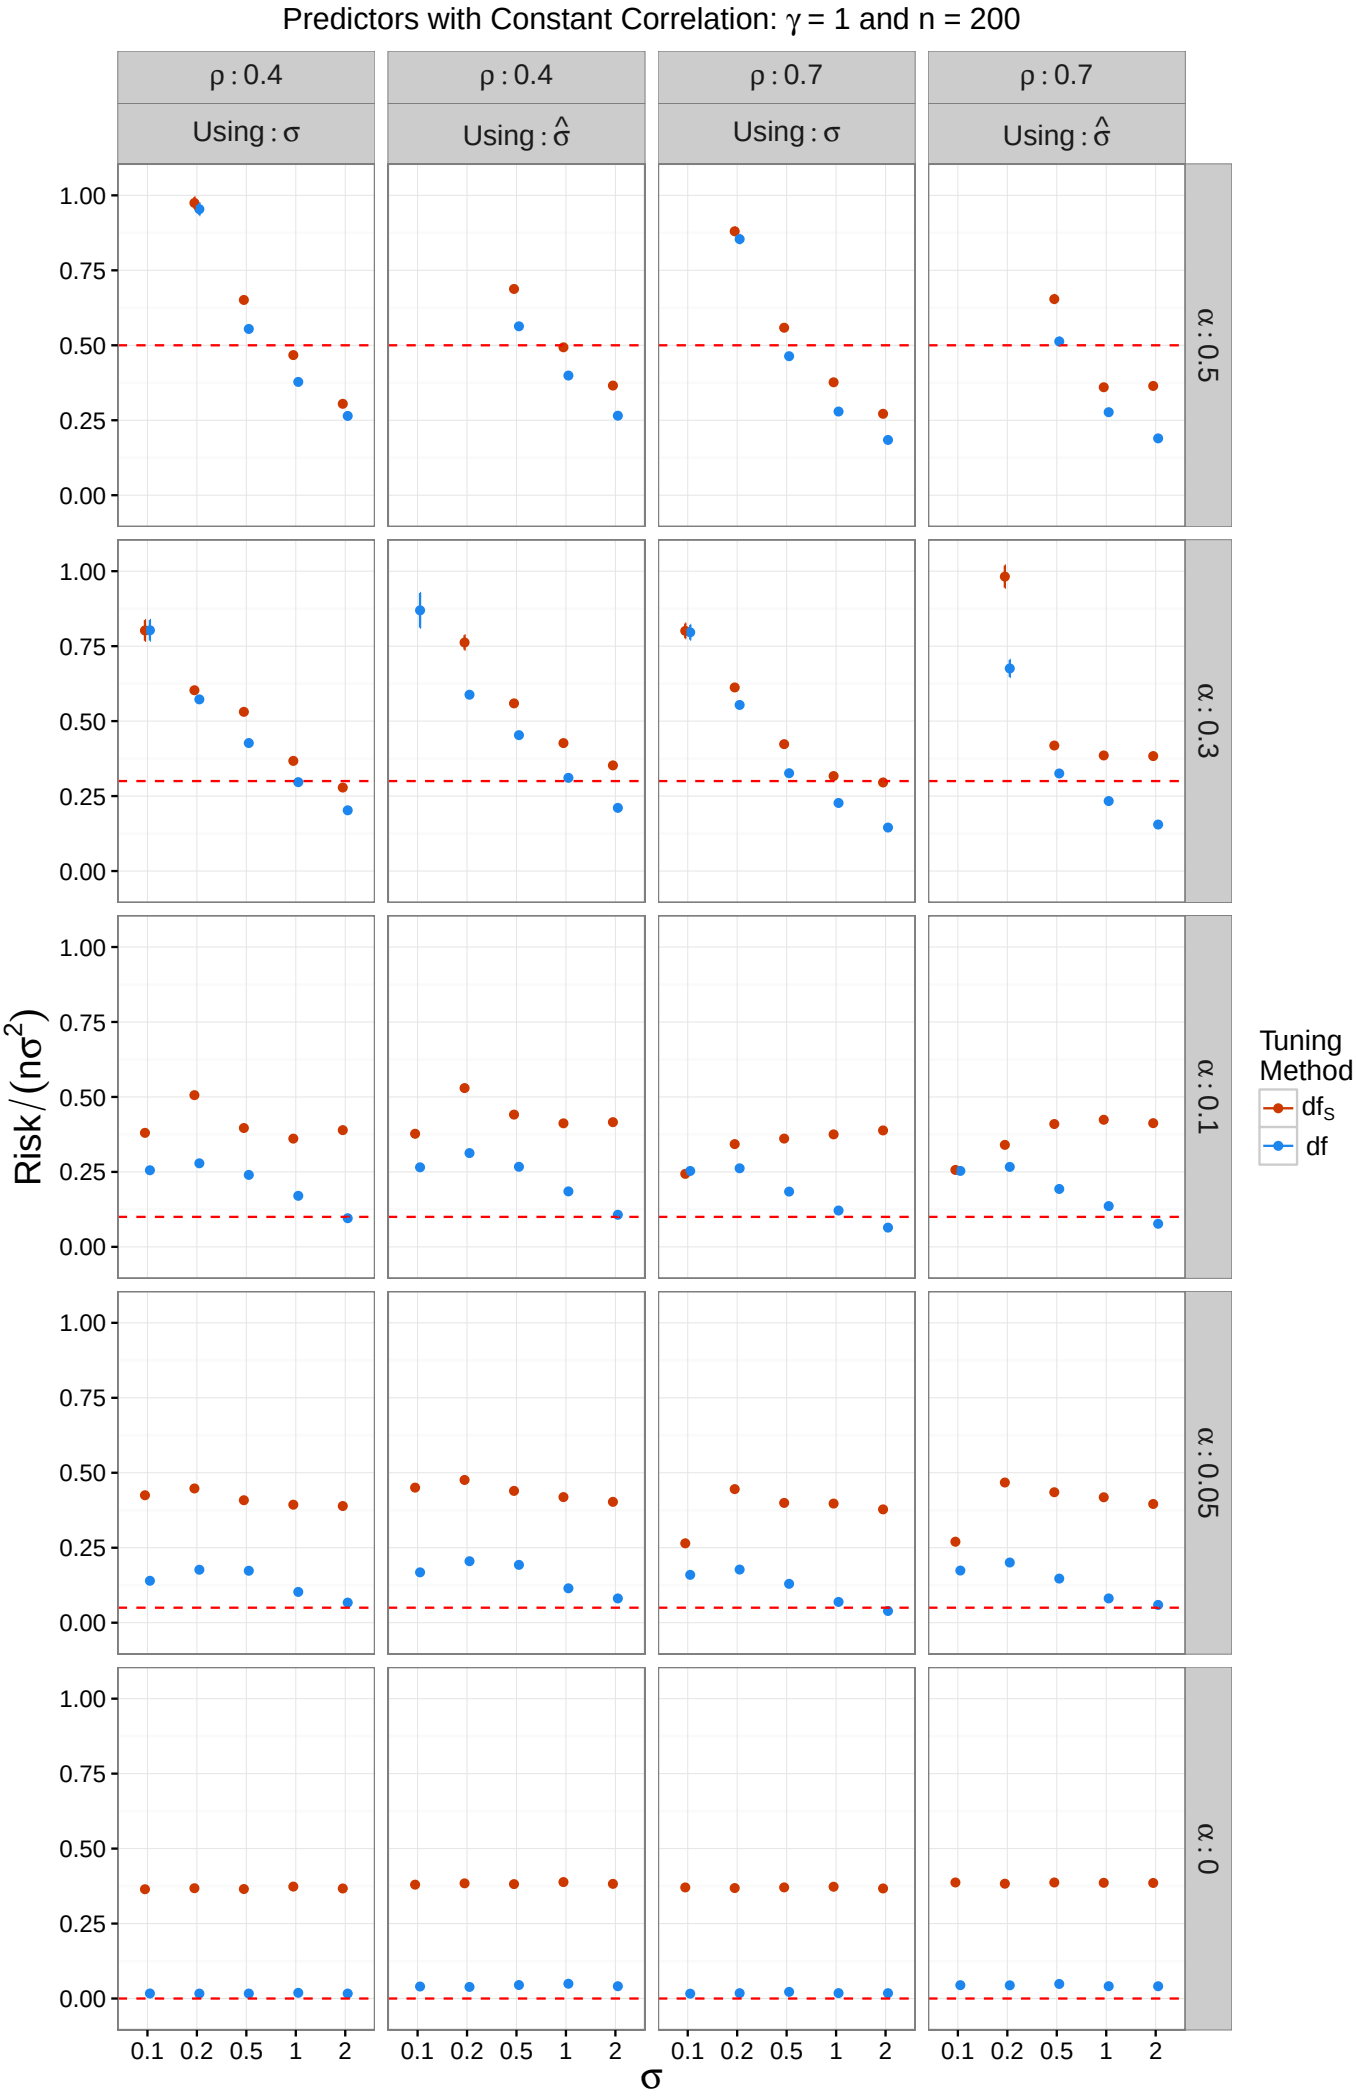
Predictors with Constant Correlation: γ = 1 and n = 200
ρ : 0.4
Using : σ
ρ : 0.4
Using : σ̂
ρ : 0.7
Using : σ
ρ : 0.7
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1
0.2
0.5
1
2
σ
Tuning Method
df_S
df

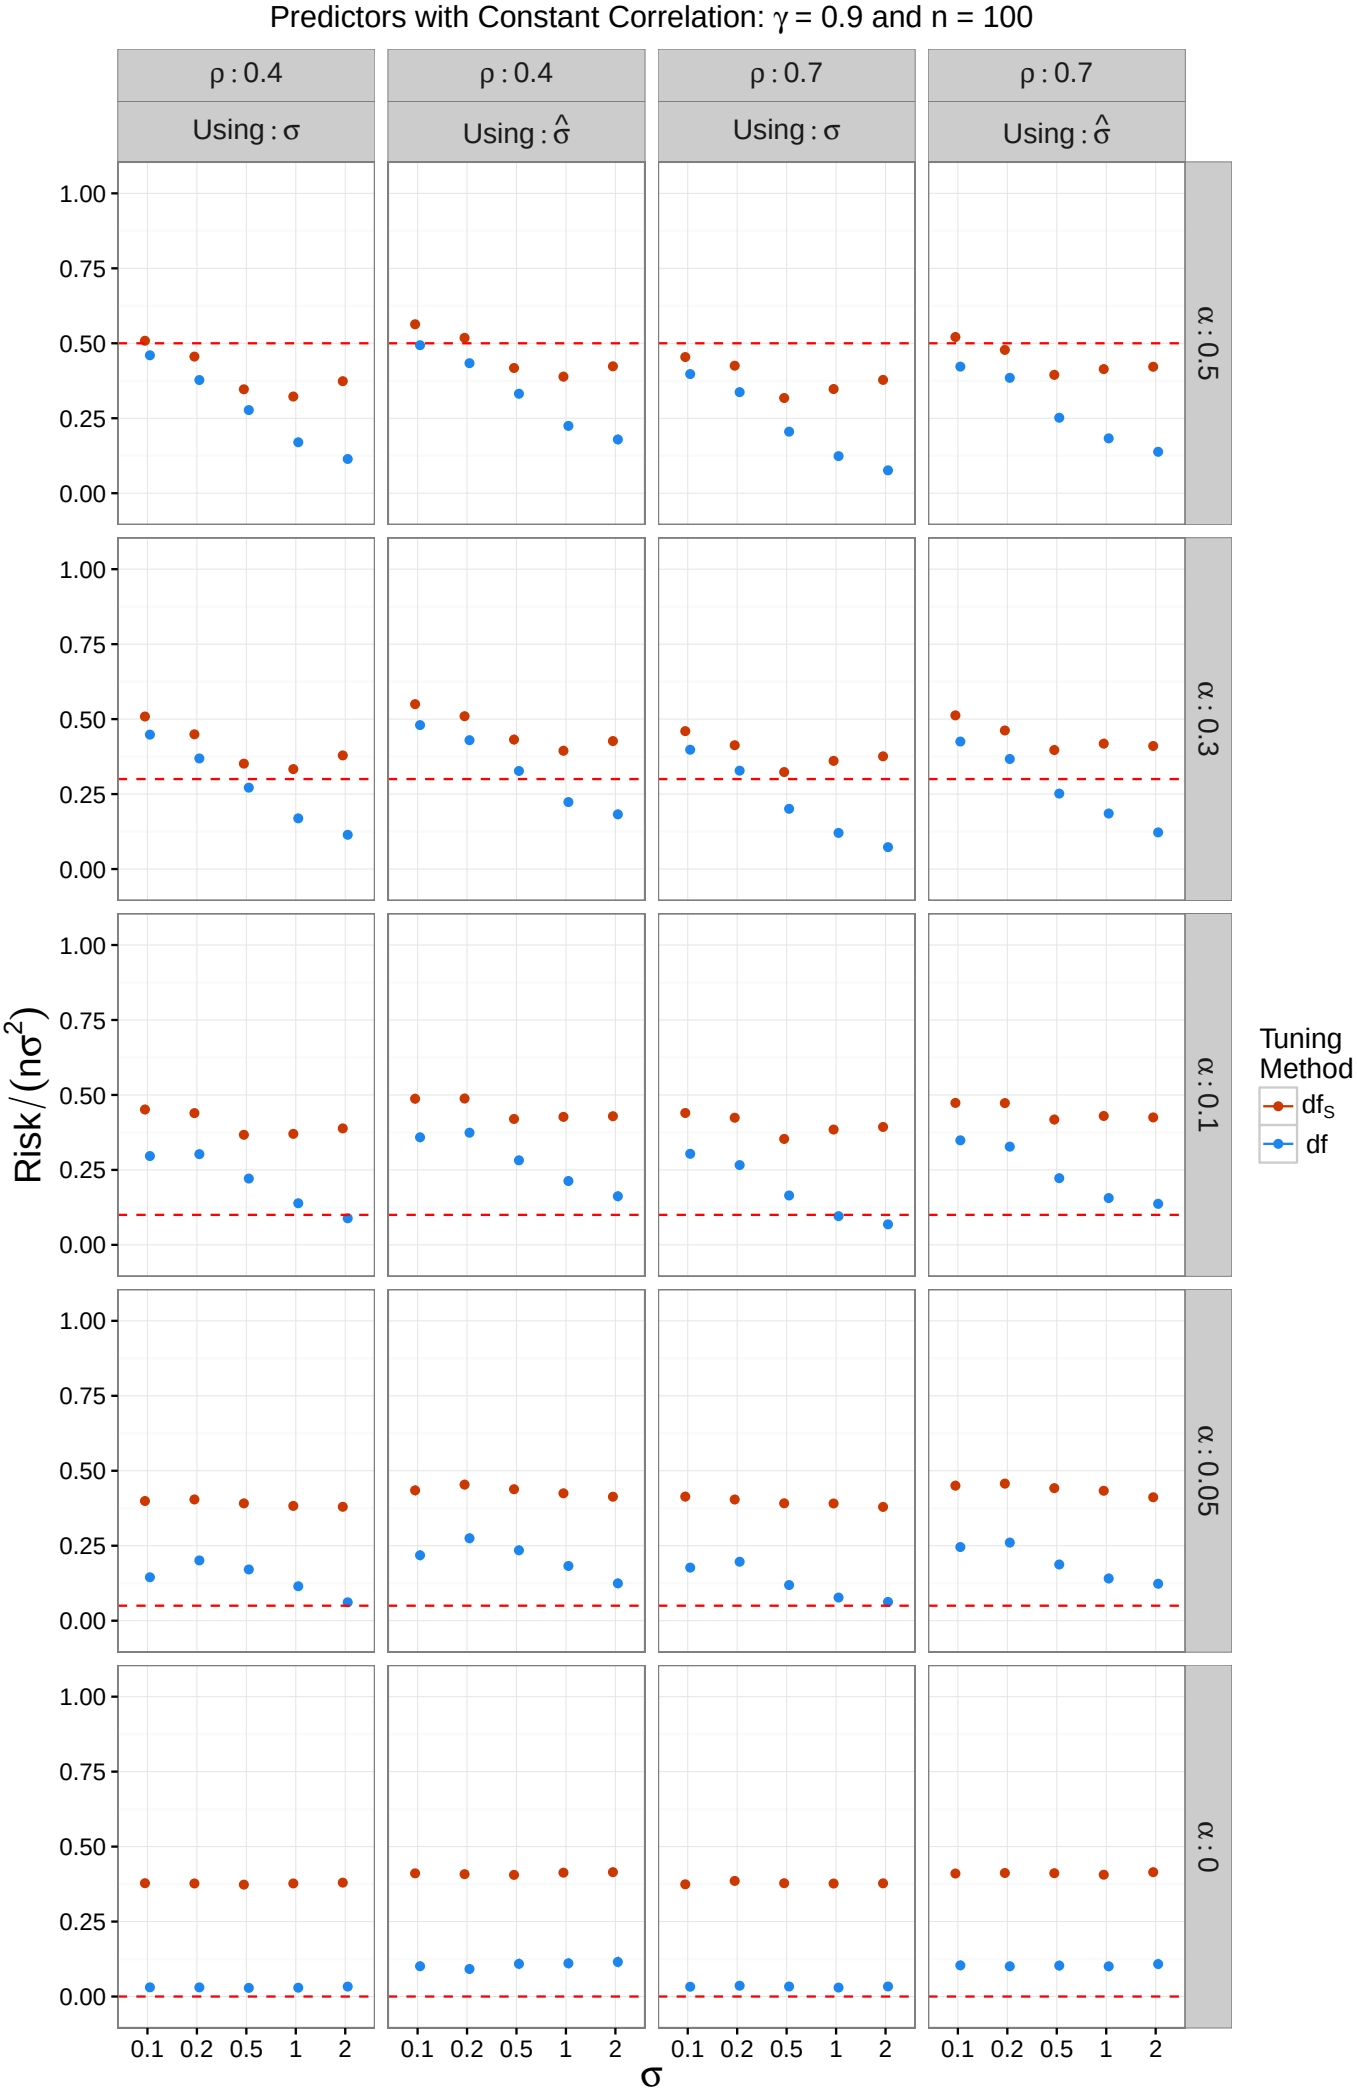
Predictors with Constant Correlation: γ = 0.9 and n = 100
ρ : 0.4
Using : σ
ρ : 0.4
Using : σ̂
ρ : 0.7
Using : σ
ρ : 0.7
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1 0.2 0.5 1 2
σ
Tuning Method
df_S
df

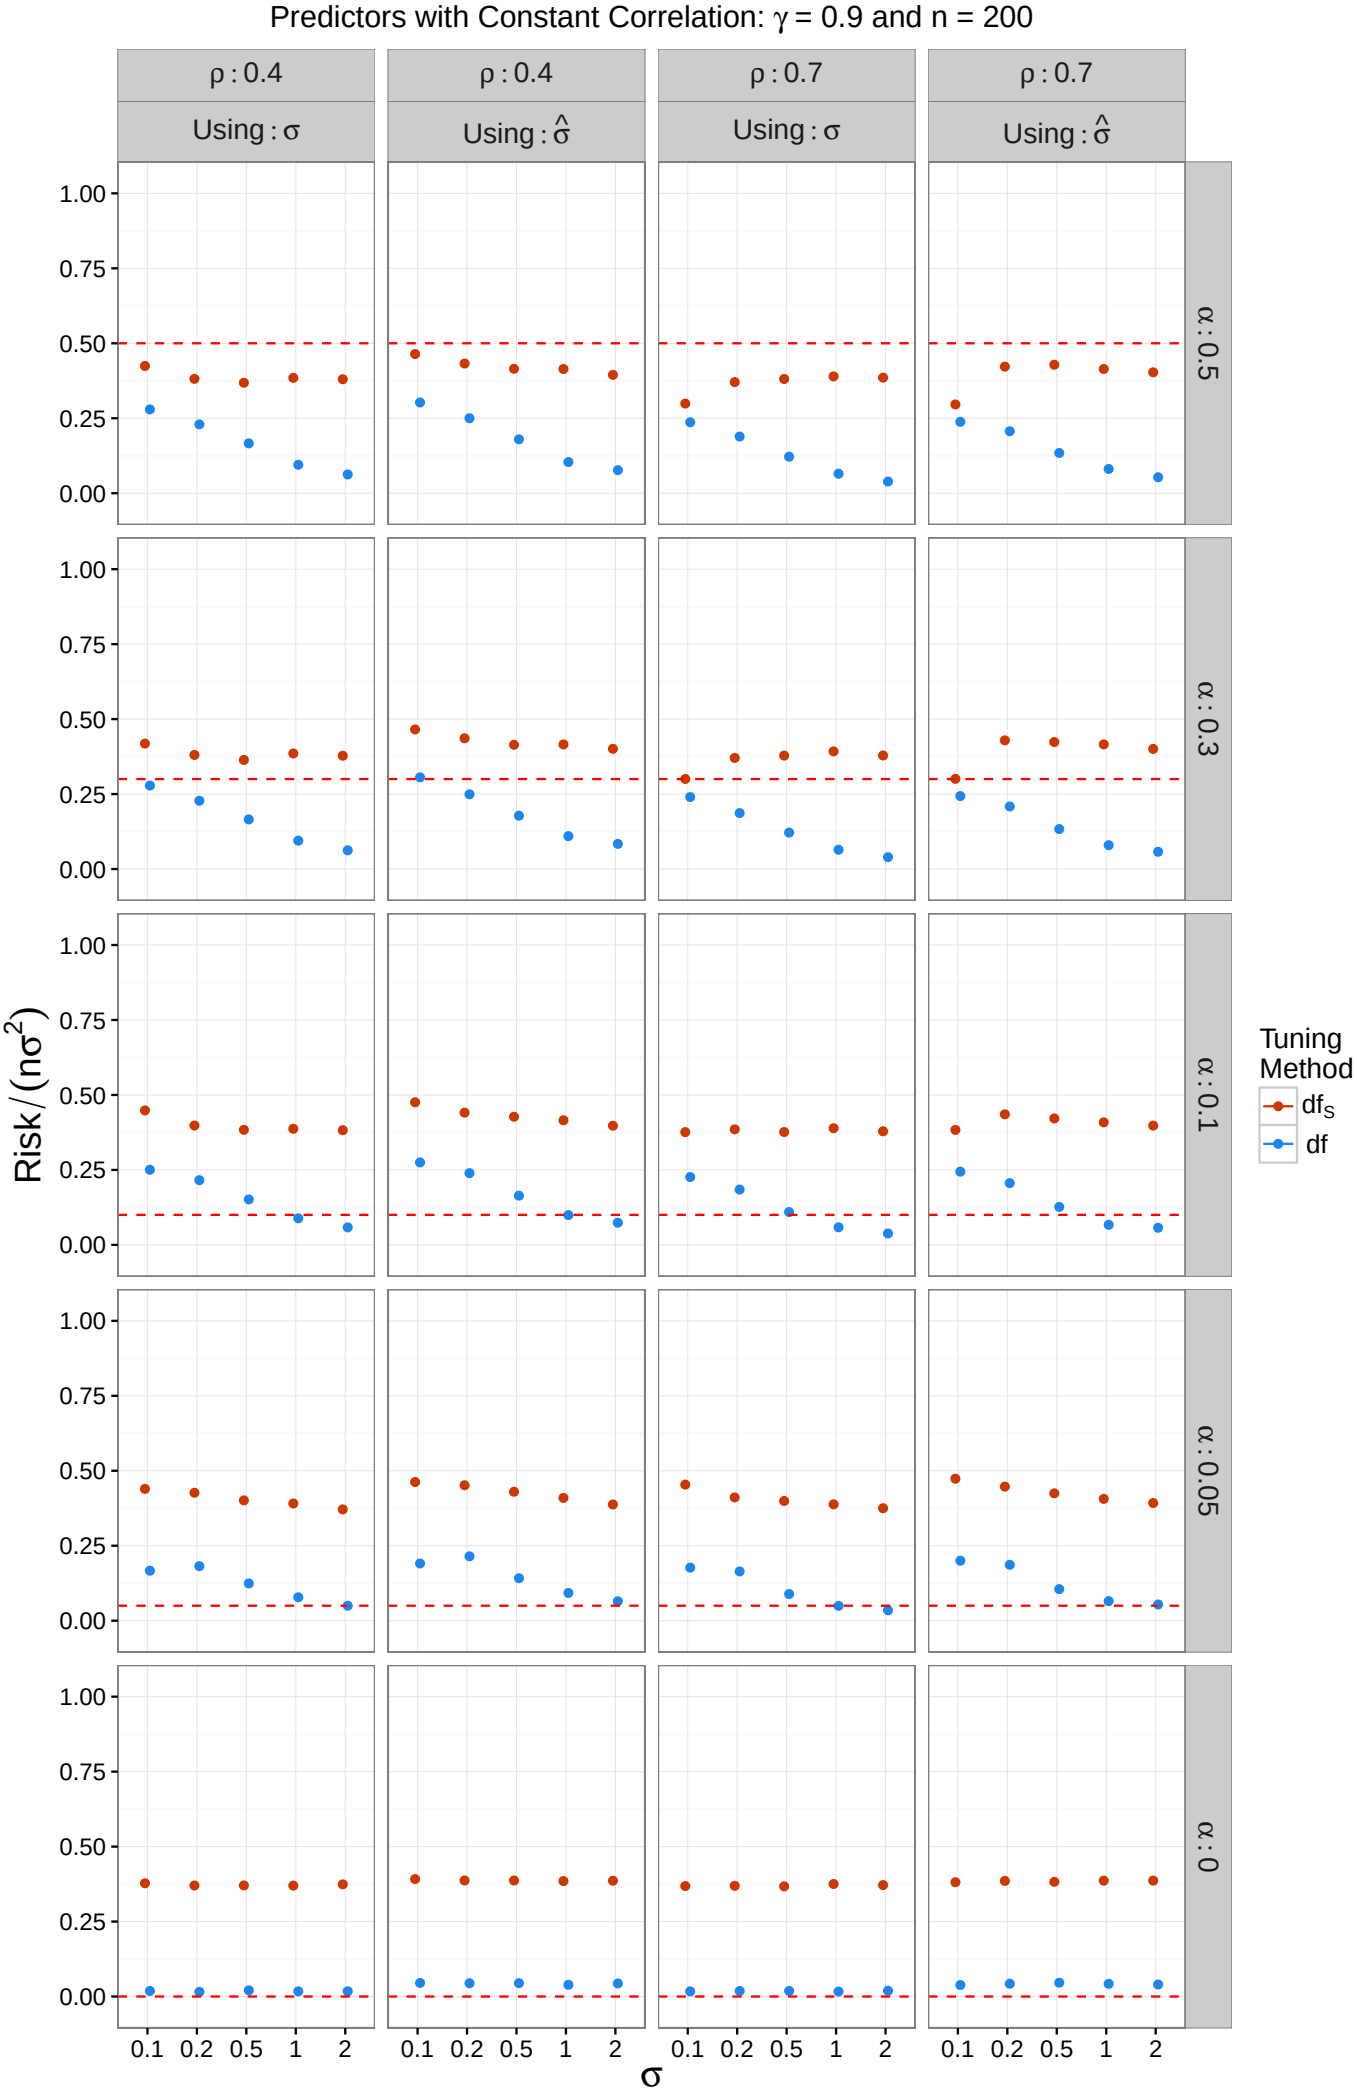
Predictors with Constant Correlation: γ = 0.9 and n = 200
ρ : 0.4
Using : σ
ρ : 0.4
Using : σ̂
ρ : 0.7
Using : σ
ρ : 0.7
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1 0.2 0.5 1 2
σ
Tuning Method
df_S
df

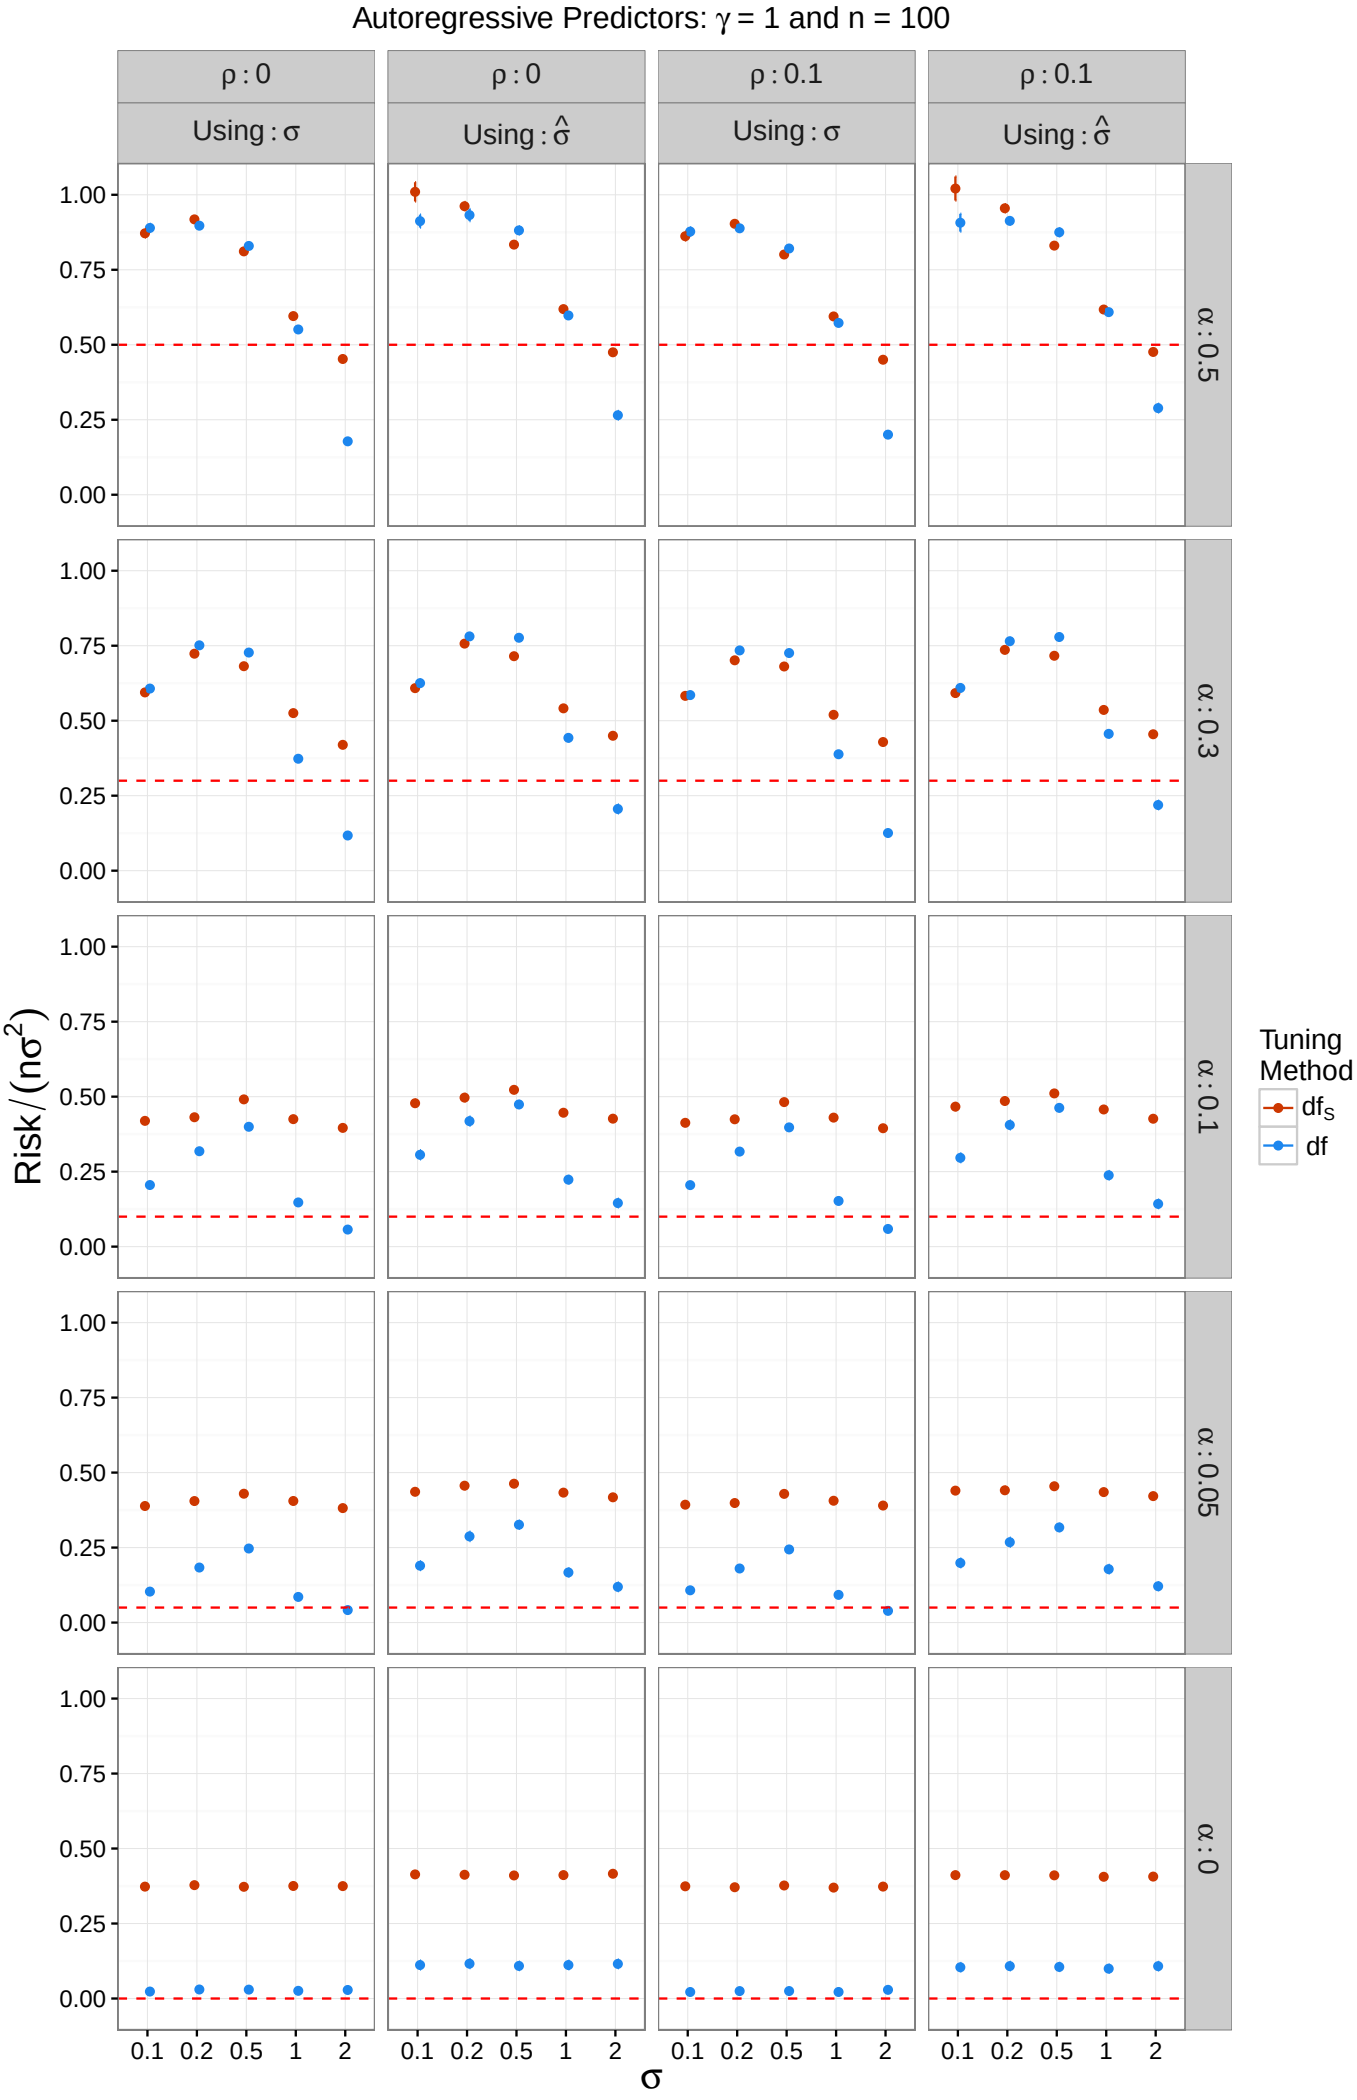
Autoregressive Predictors: γ = 1 and n = 100
ρ : 0
Using : σ
ρ : 0
Using : σ̂
ρ : 0.1
Using : σ
ρ : 0.1
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1
0.2
0.5
1
2
σ
Tuning Method
df_S
df

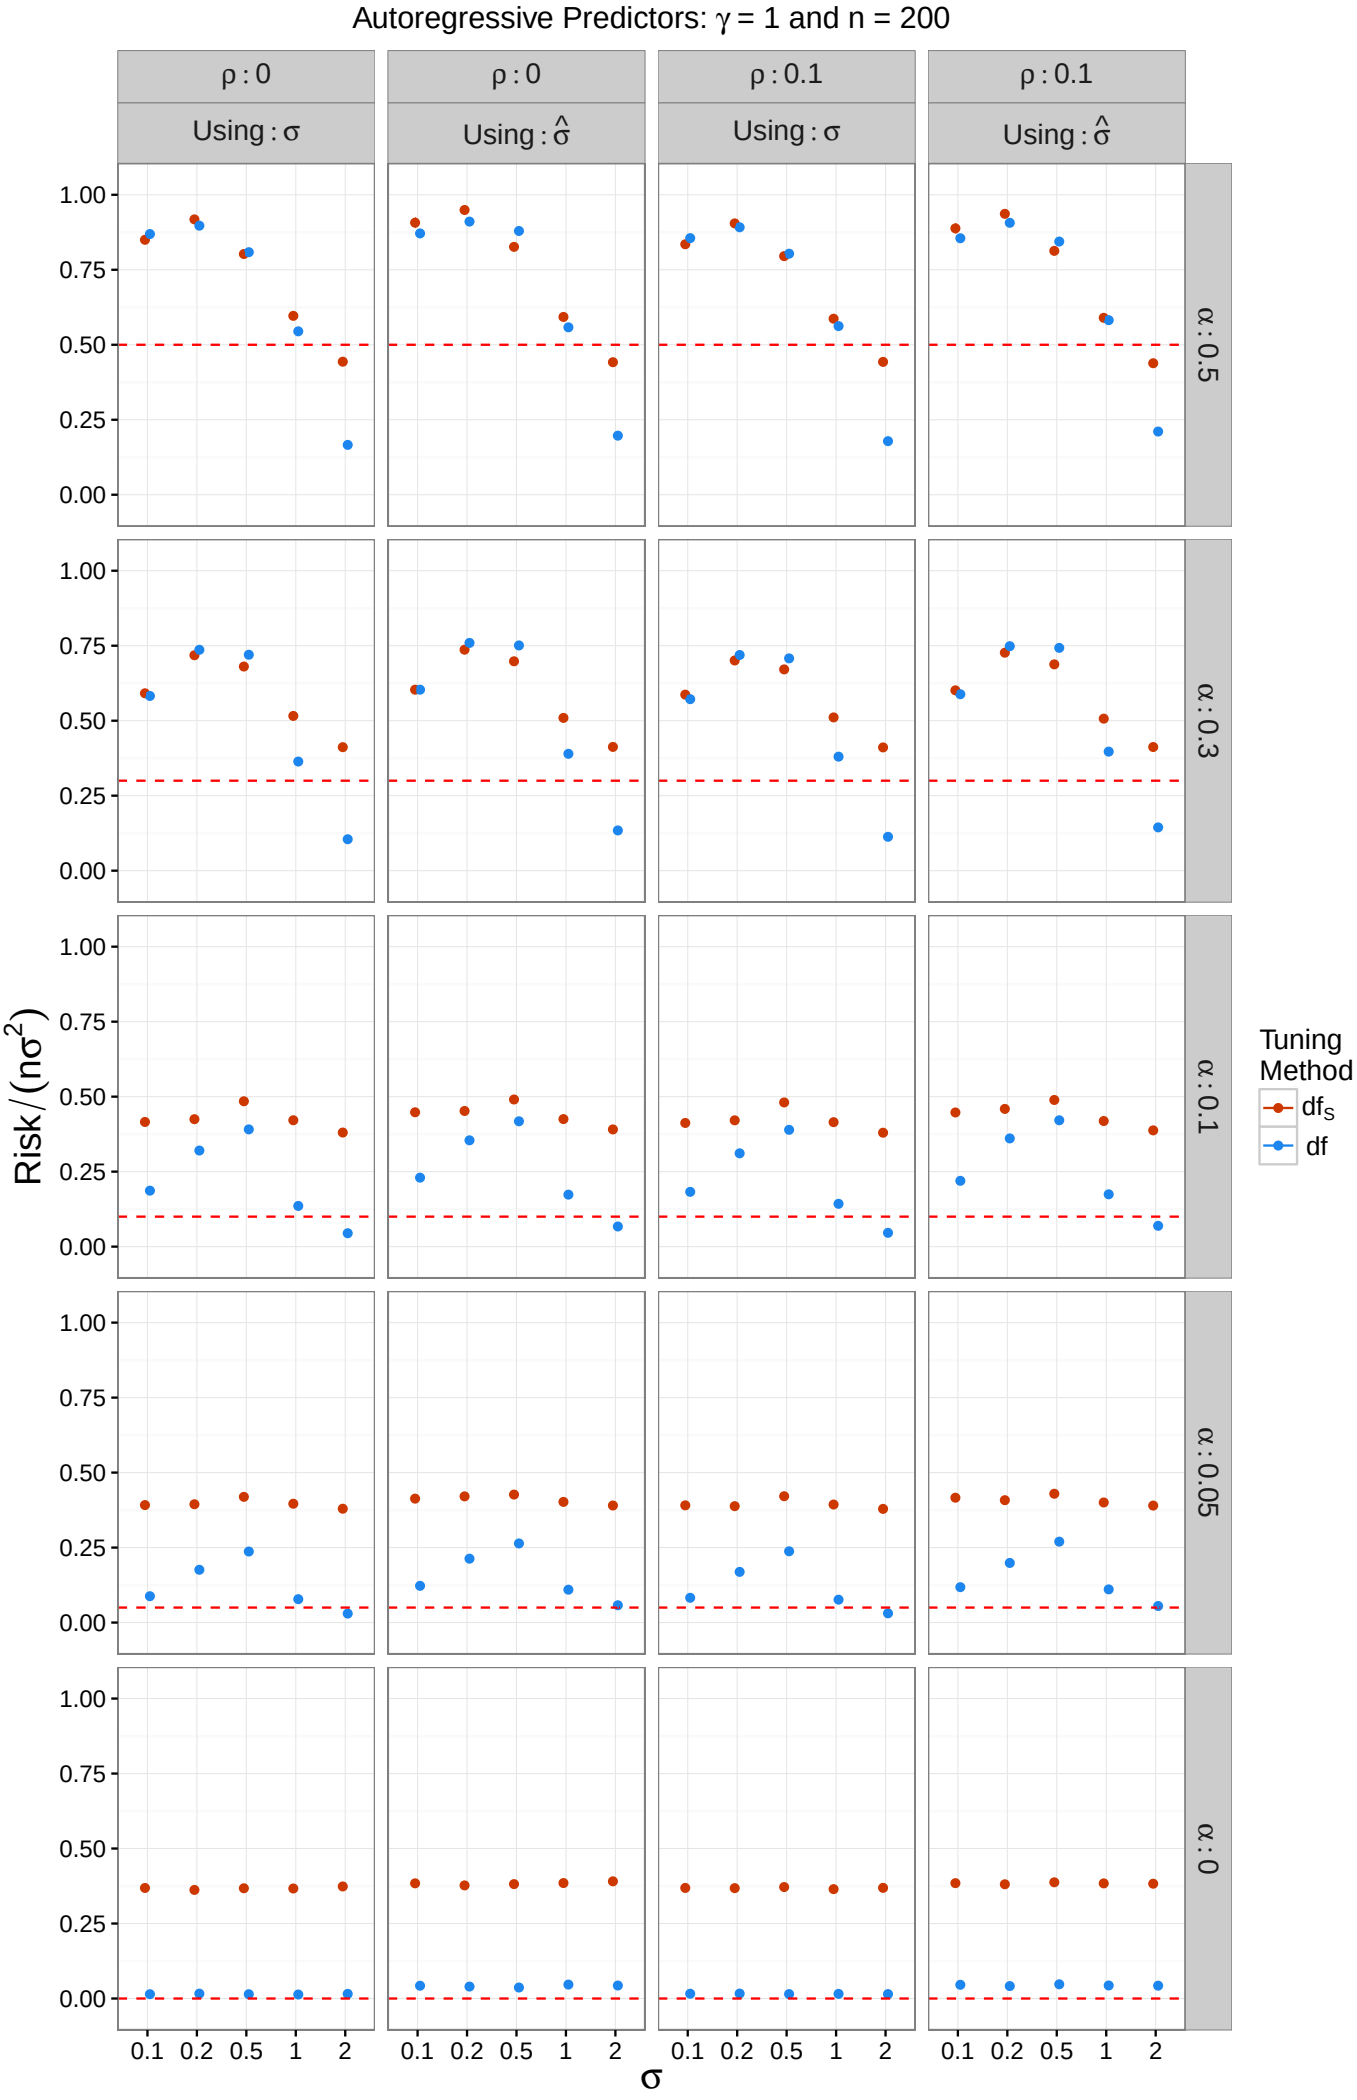
Autoregressive Predictors: γ = 1 and n = 200
ρ : 0
Using : σ
ρ : 0
Using : σ̂
ρ : 0.1
Using : σ
ρ : 0.1
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1
0.2
0.5
1
2
σ
Tuning Method
df_S
df

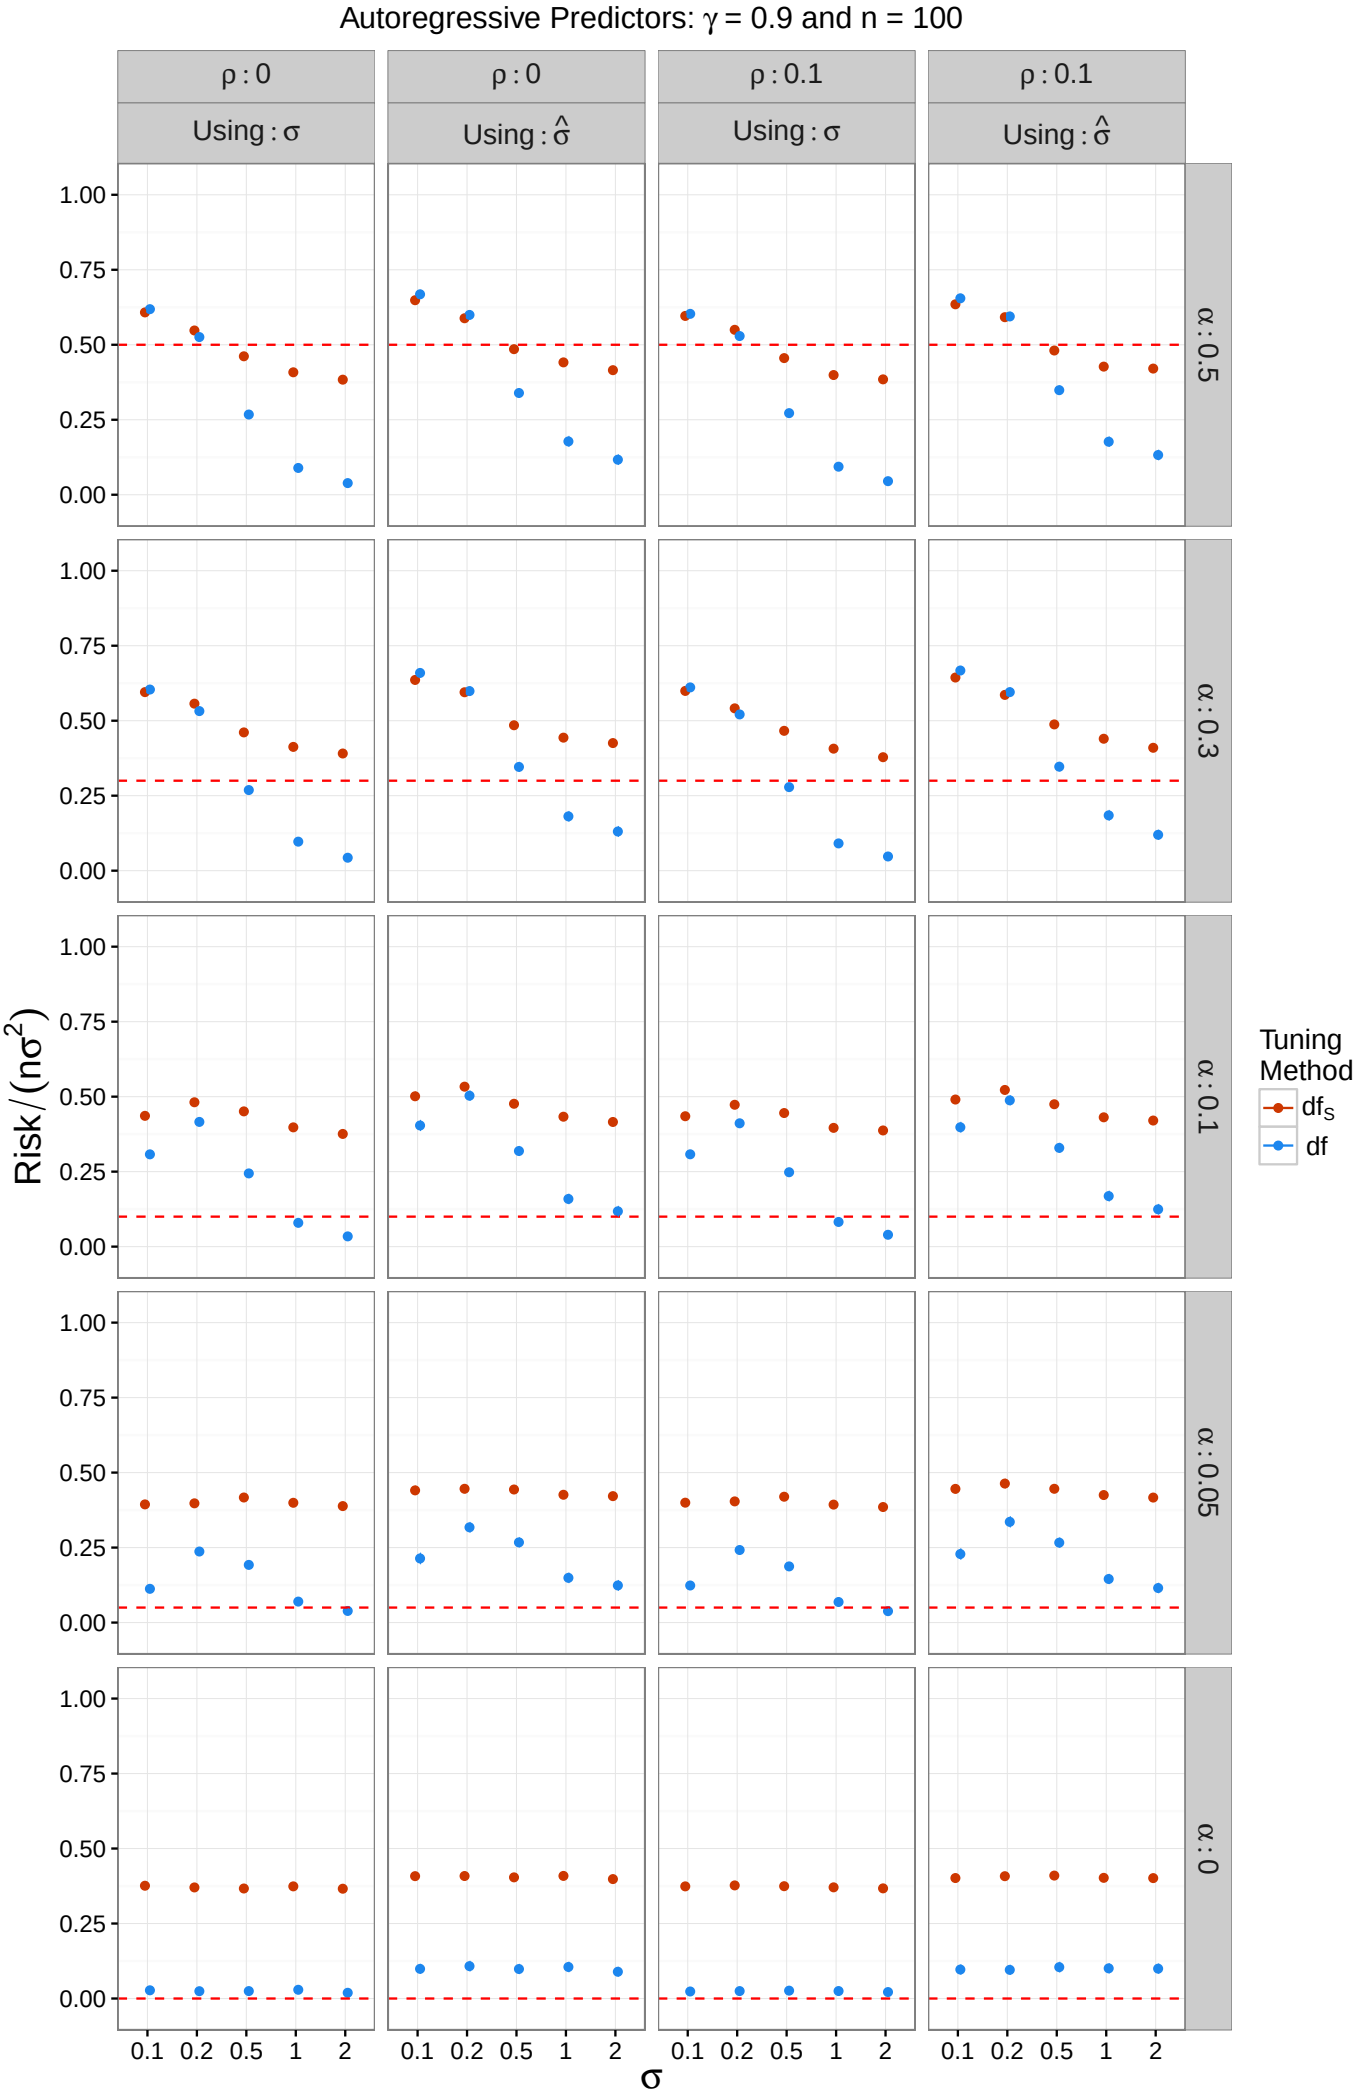
Autoregressive Predictors: γ = 0.9 and n = 100
ρ : 0
Using : σ
ρ : 0
Using : σ̂
ρ : 0.1
Using : σ
ρ : 0.1
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1 0.2 0.5 1 2
σ
Tuning Method
df_S
df

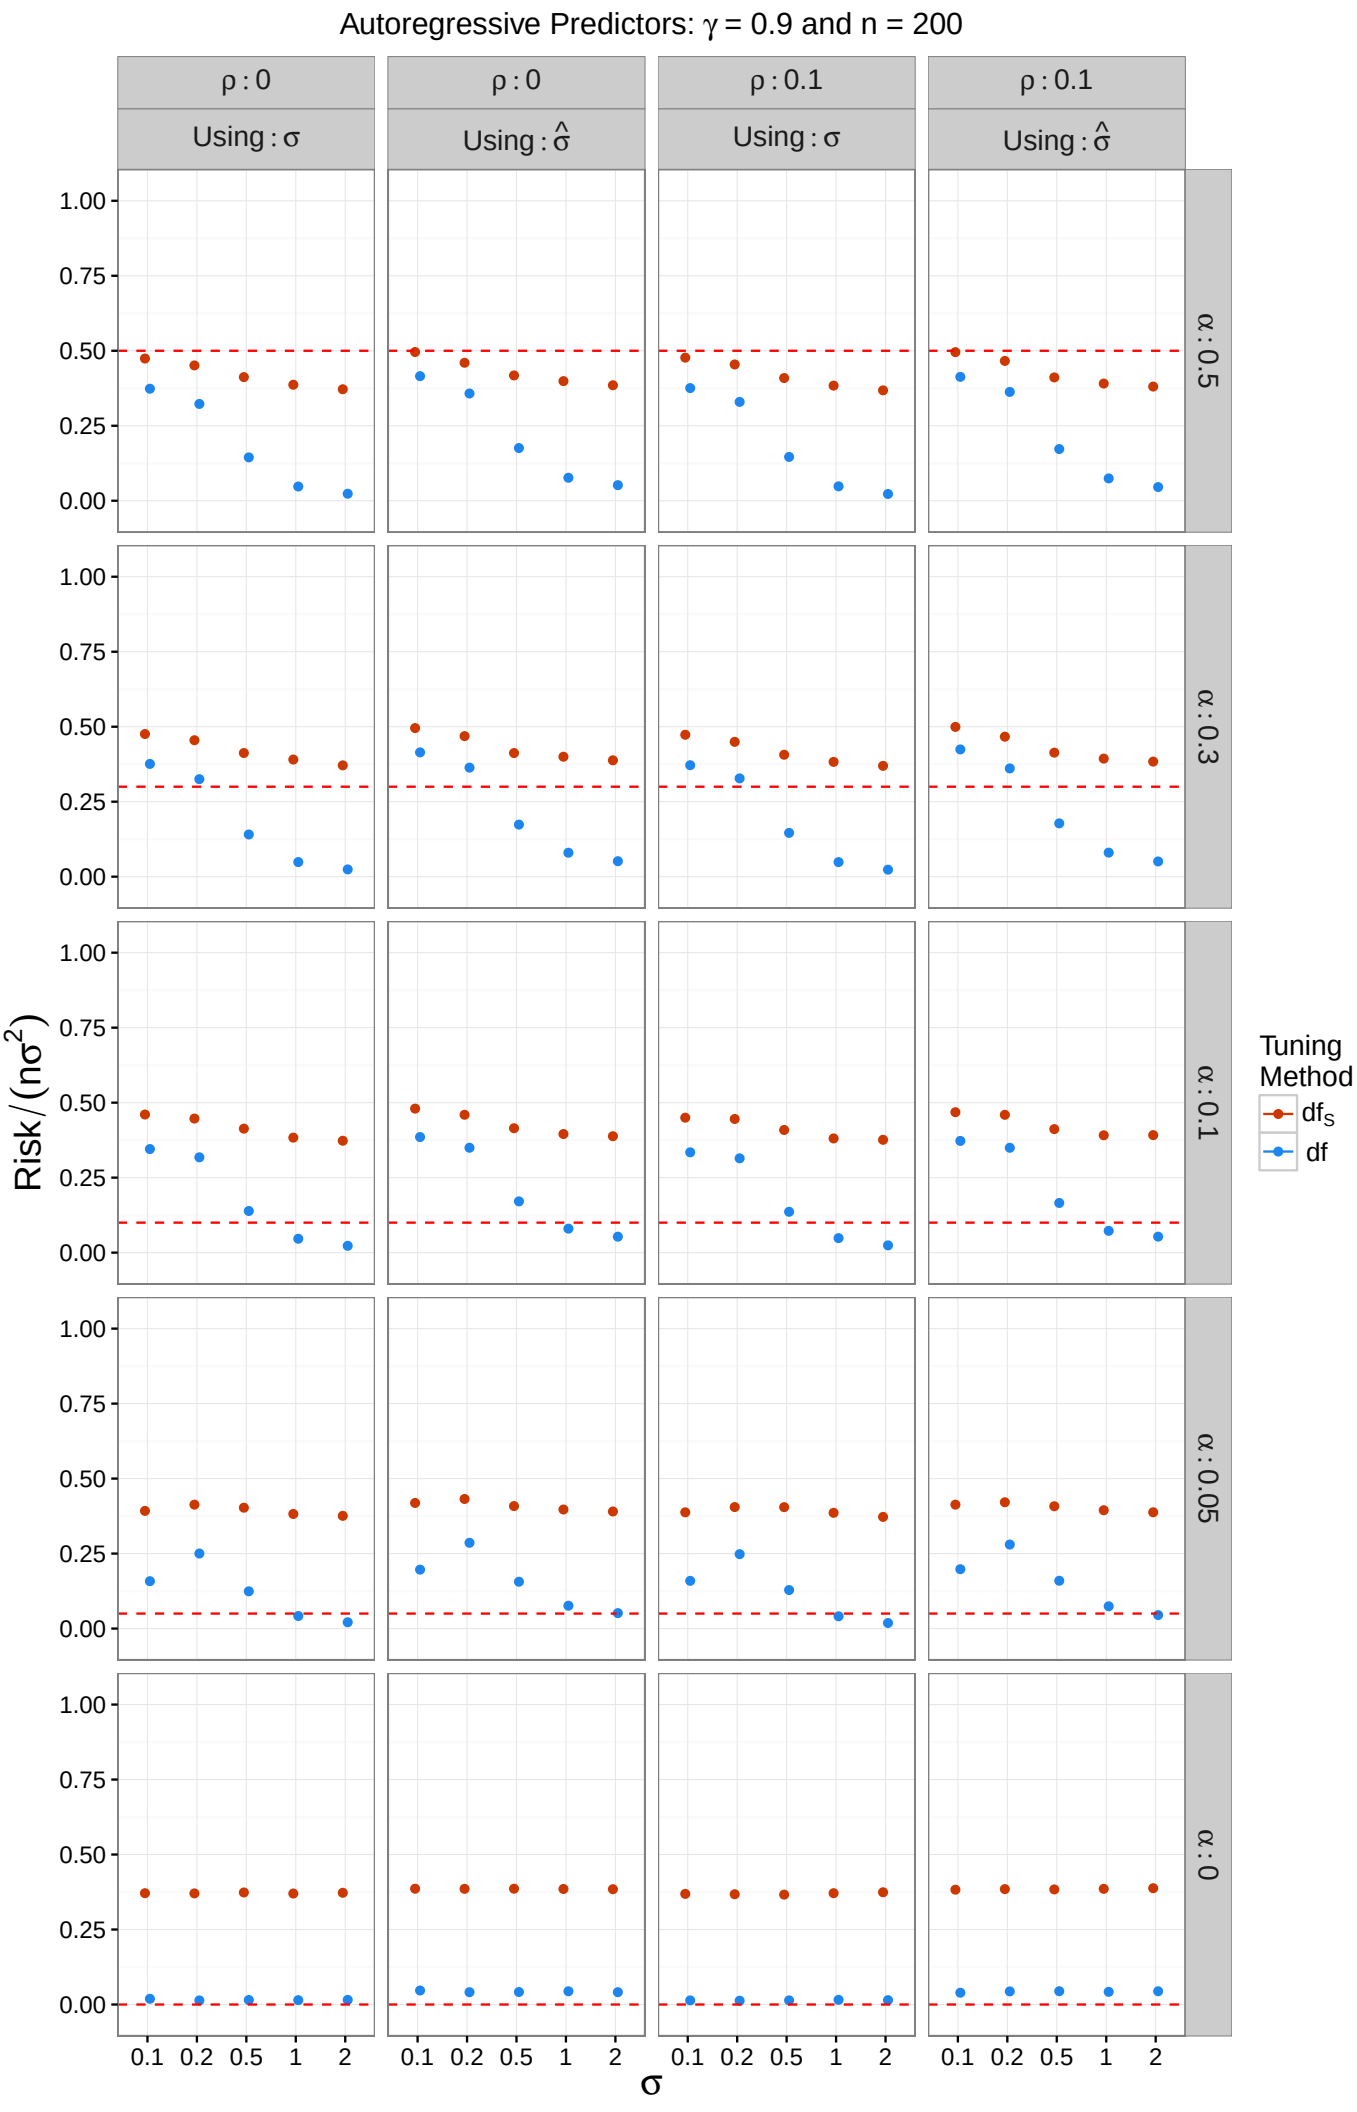
Autoregressive Predictors: γ = 0.9 and n = 200
ρ : 0
Using : σ
ρ : 0
Using : σ̂
ρ : 0.1
Using : σ
ρ : 0.1
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1 0.2 0.5 1 2
σ
Tuning Method
df_S
df

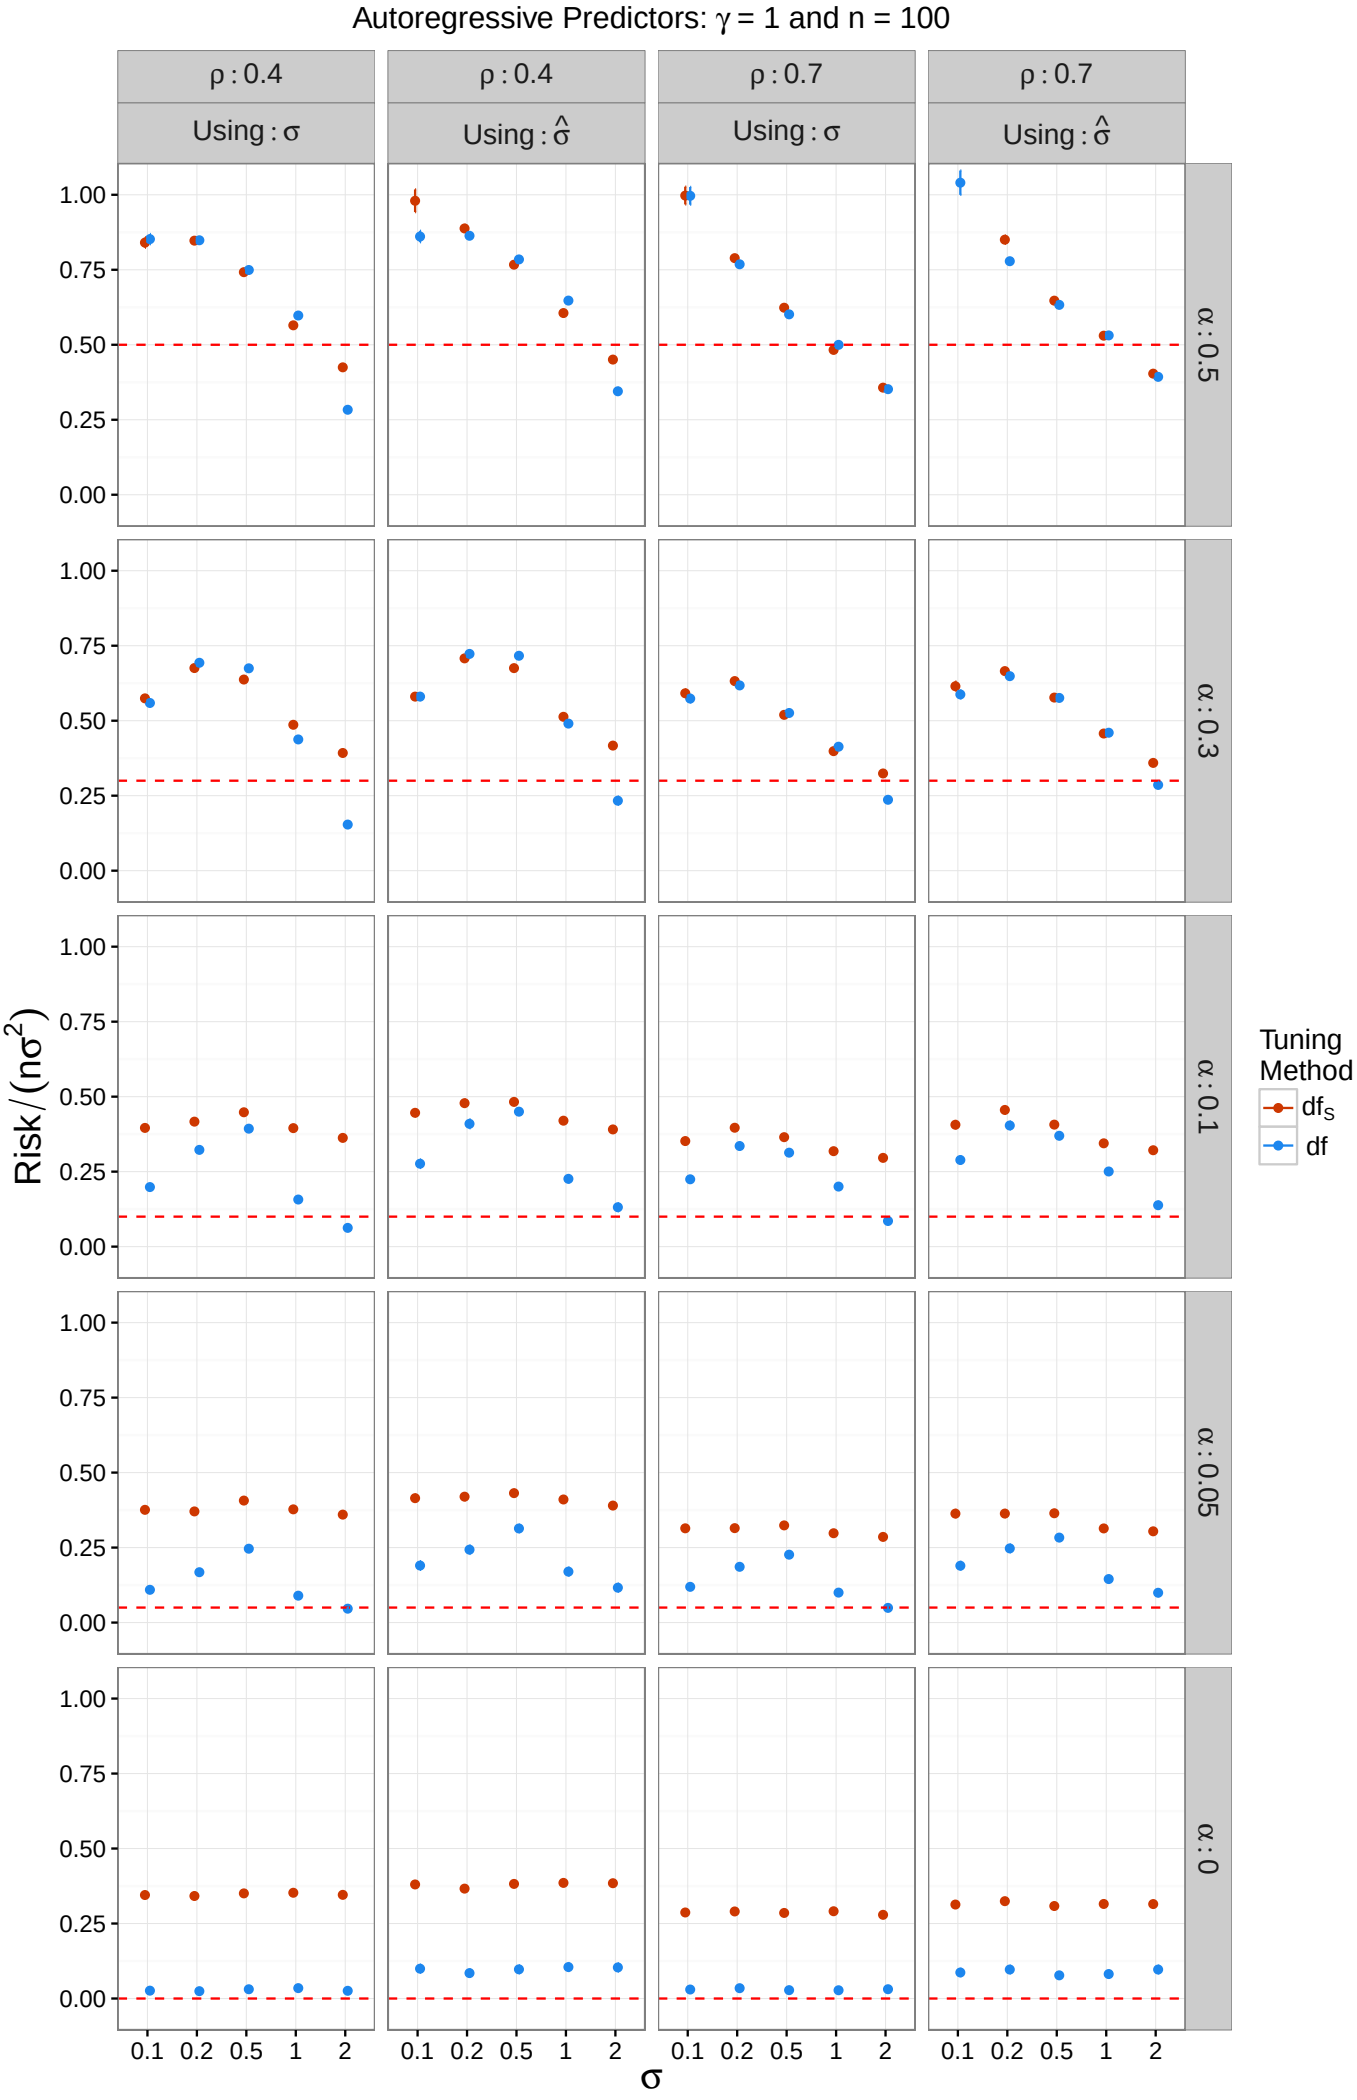
Autoregressive Predictors: γ = 1 and n = 100
ρ : 0.4
Using : σ
ρ : 0.4
Using : σ̂
ρ : 0.7
Using : σ
ρ : 0.7
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1
0.2
0.5
1
2
σ
Tuning Method
df_S
df

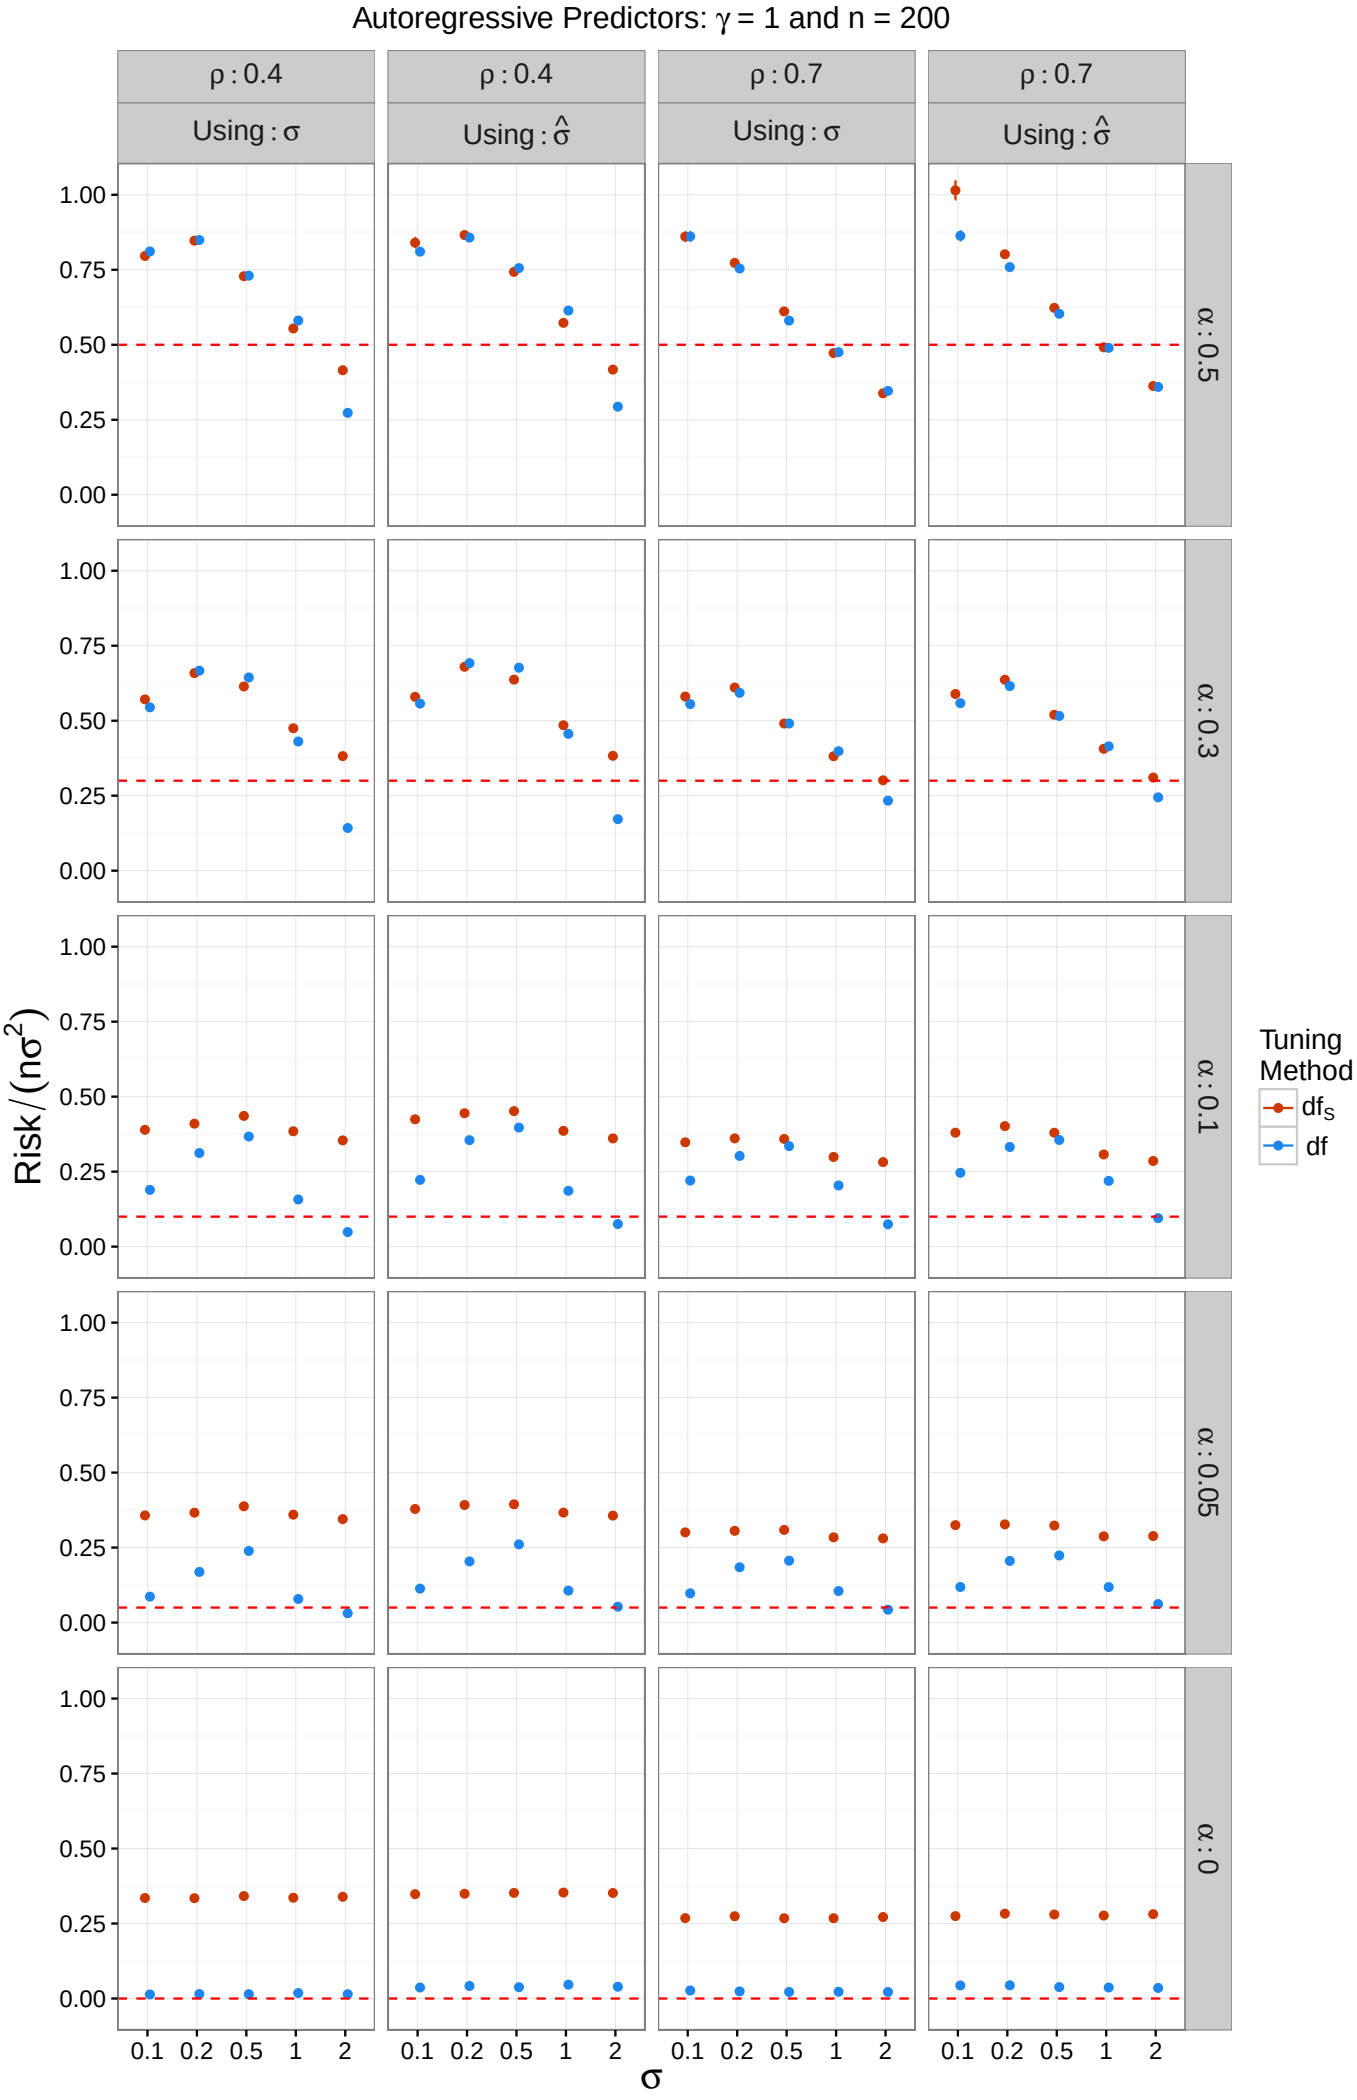
Autoregressive Predictors: γ = 1 and n = 200
ρ : 0.4
Using : σ
ρ : 0.4
Using : σ̂
ρ : 0.7
Using : σ
ρ : 0.7
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1 0.2 0.5 1 2
σ
Tuning Method
df_S
df

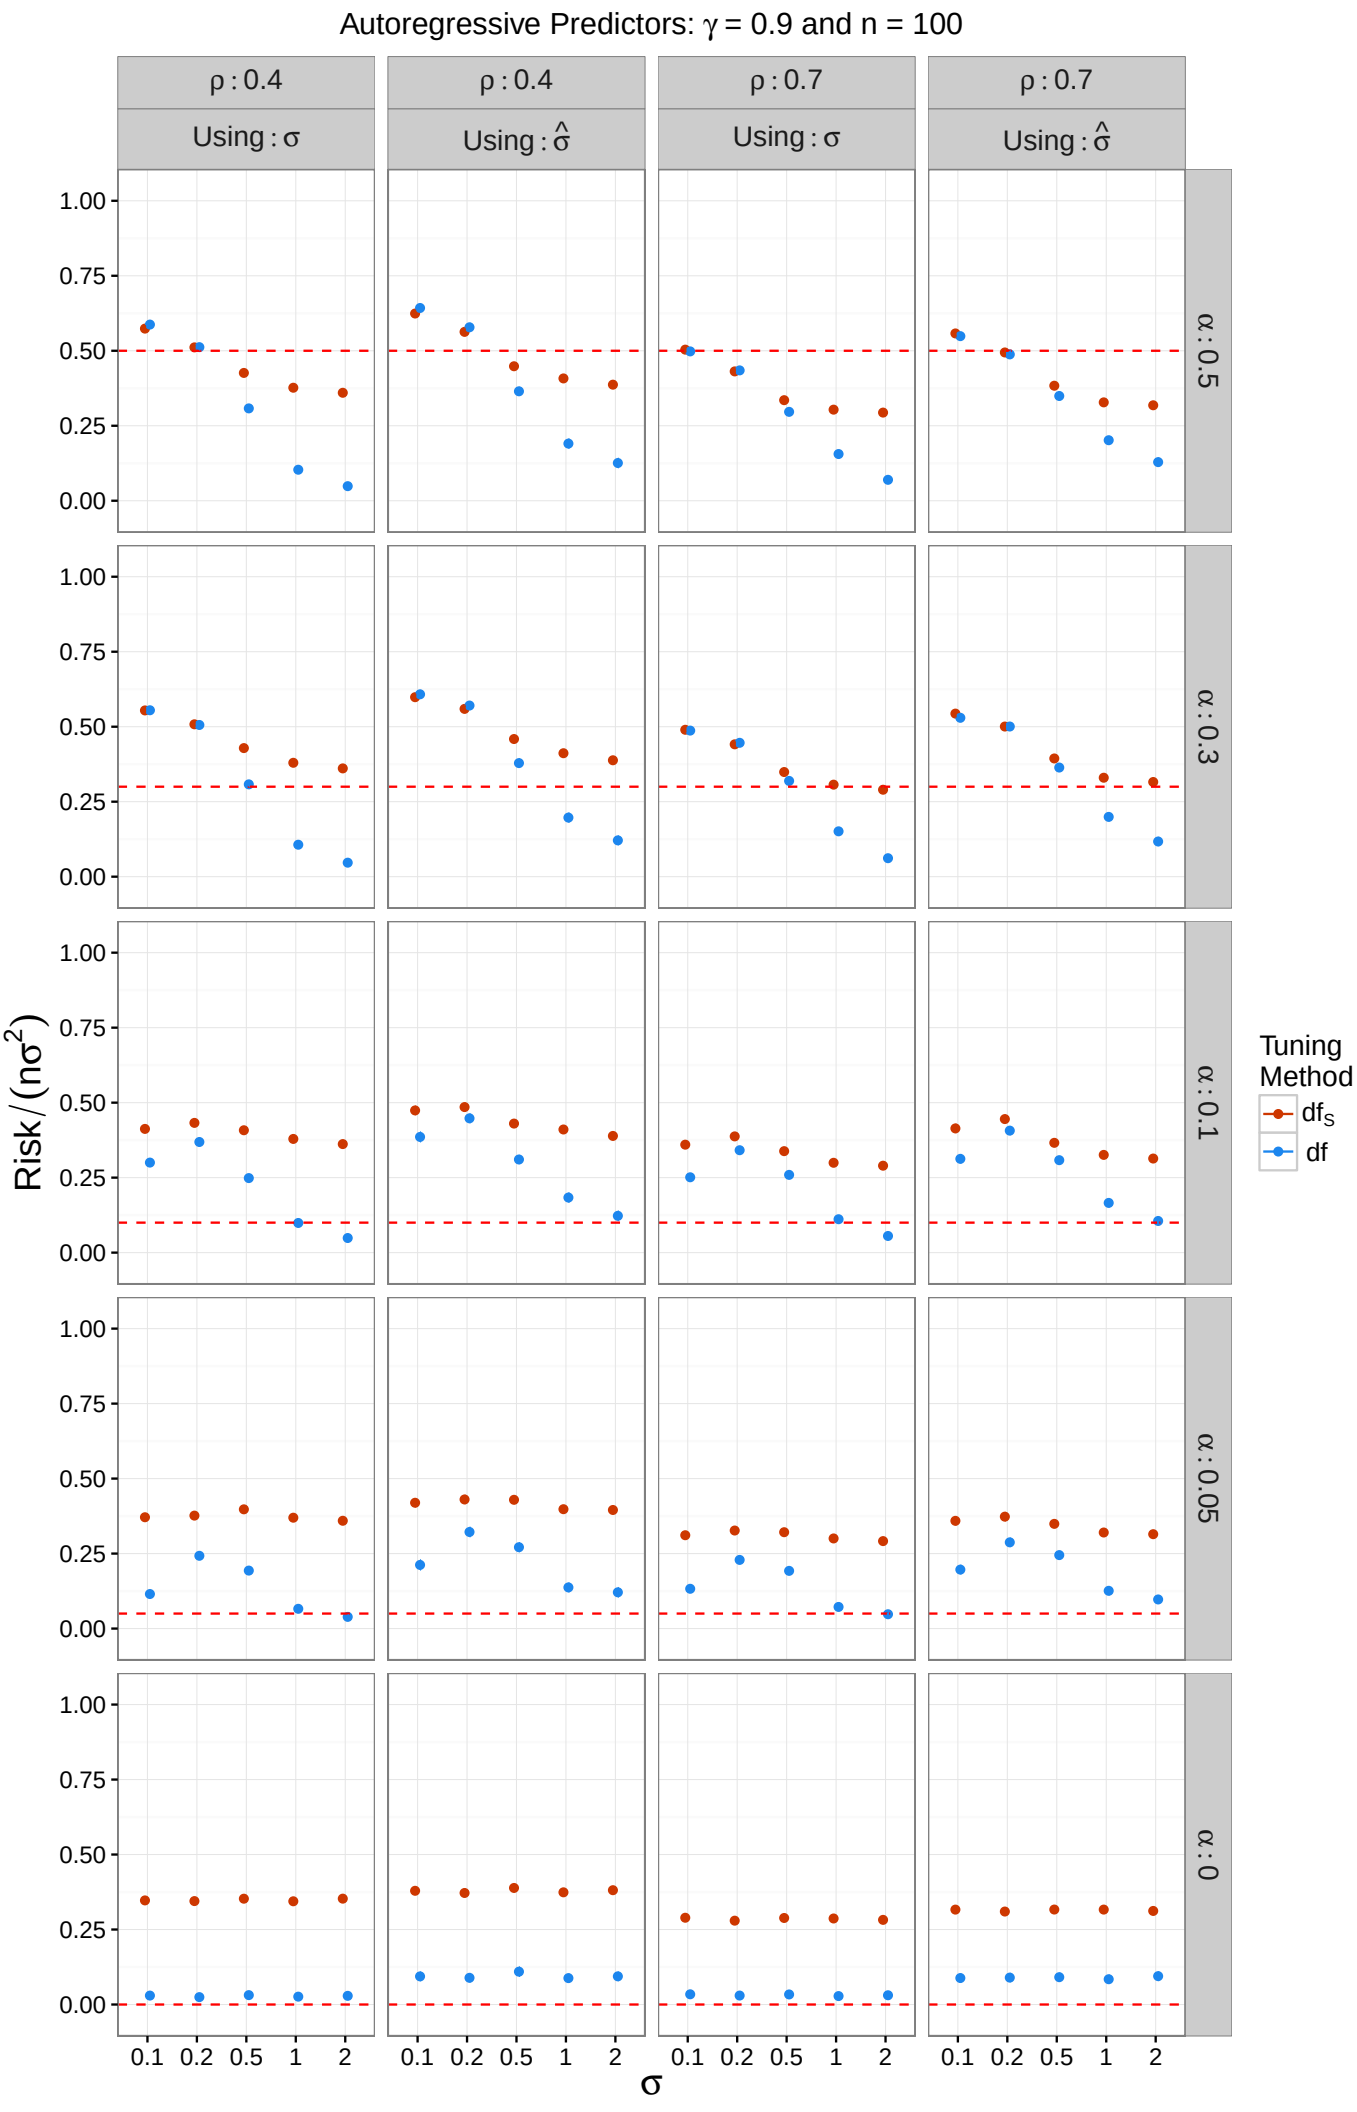
Autoregressive Predictors: γ = 0.9 and n = 100
ρ : 0.4
Using : σ
ρ : 0.4
Using : σ̂
ρ : 0.7
Using : σ
ρ : 0.7
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1 0.2 0.5 1 2
σ
Tuning Method
df_S
df

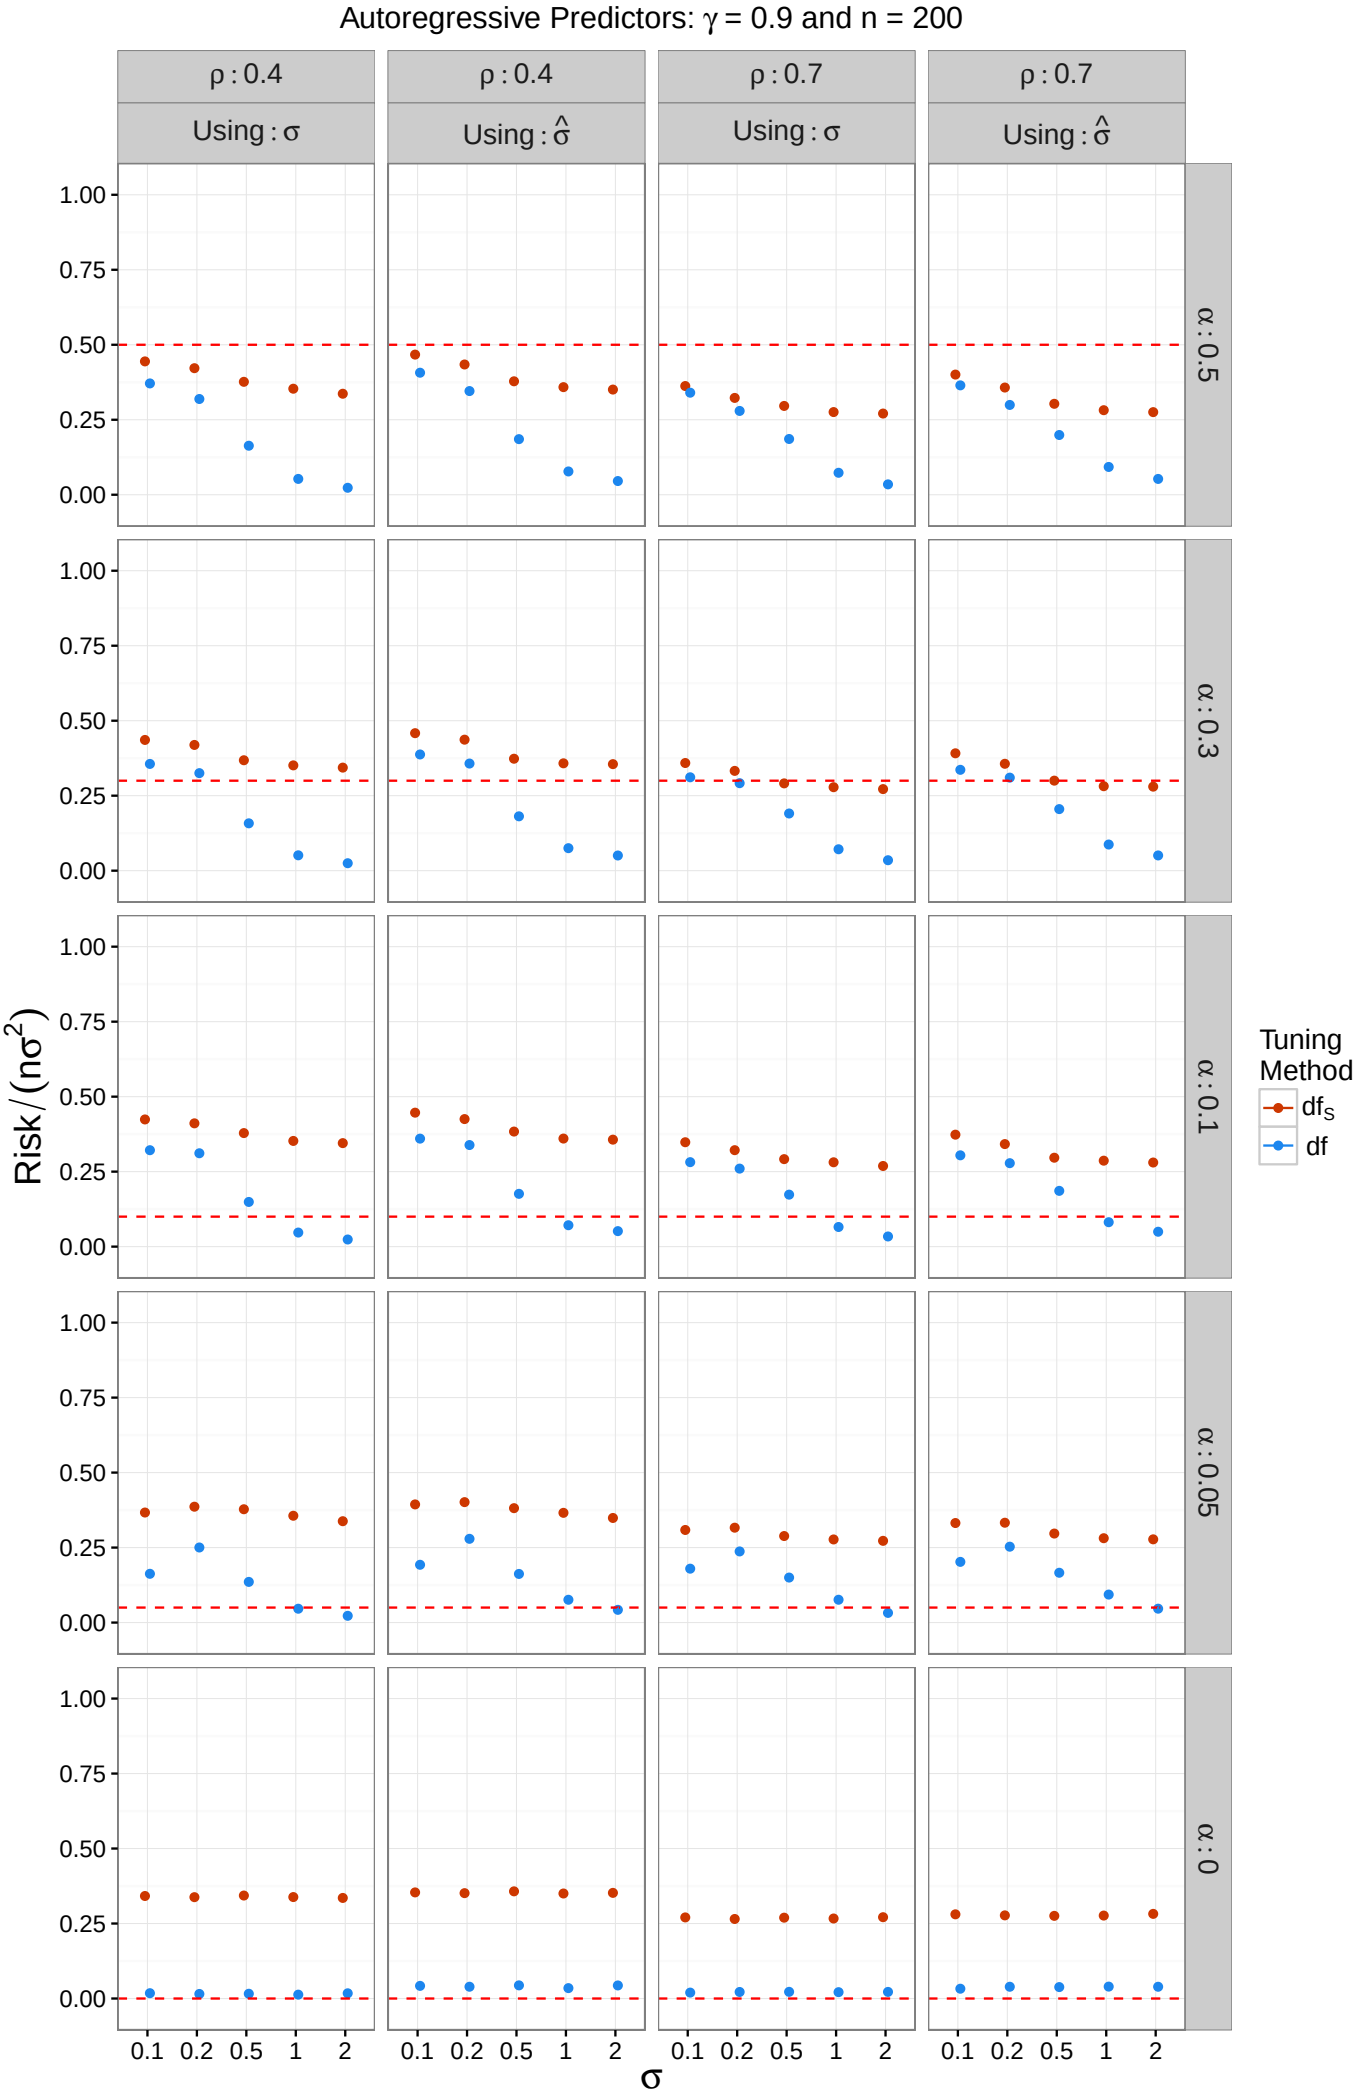
Autoregressive Predictors: γ = 0.9 and n = 200
ρ : 0.4
Using : σ
ρ : 0.4
Using : σ̂
ρ : 0.7
Using : σ
ρ : 0.7
Using : σ̂
α : 0.5
α : 0.3
α : 0.1
α : 0.05
α : 0
Risk/(nσ²)
1.00
0.75
0.50
0.25
0.00
0.1 0.2 0.5 1 2
σ
Tuning Method
df_S
df

Department of Mathematical Sciences, University of Copenhagen, Universitetsparken 5, 2100 Copenhagen Ø, Denmark

E-mail address, Corresponding author: frm@math.ku.dk

E-mail address: Niels.R.Hansen@math.ku.dk